

BIG.D

Interdisciplinary One-Shot Learning: Mimicking Human Cognition for Data-Efficient Innovation Across Design, Engineering, Business, and Culture

1st Shuvadra Banik *
BGMEA University of Fashion &
Technology
Dhaka, Bangladesh
Shuvadrabanik@outlook.com

2nd Zannatul Maua Tuly
BGMEA University of Fashion &
Technology
Dhaka, Bangladesh
ZannatulMT@proton.me

3rd Md Fuad Hasan
Sichuan University
Sichuan, China
MdFuadH@proton.me

Abstract—Machine learning has made great strides, but its progress is often held back by a need for massive datasets. This is a major hurdle in fields where data is scarce or expensive. Humans, on the other hand, can learn from just a few examples. This paper introduces a new one-shot learning framework that mimics this human ability. We built a model that learns similarity in a way that's similar to how people do it, by looking at the general appearance of things. This makes our model more transparent and easier to understand than many other AI systems. We also developed a new optimization algorithm to make the learning process more efficient. We show that our framework can be used to solve problems in a variety of fields, including design, engineering, business, and culture. Our experiments show that our model performs well even with very little data, and that it is also easy to interpret. This work is a step towards building more human-like AI, and it also opens up new possibilities for innovation in many different areas.

Keywords—*Interdisciplinary One-Shot Learning, Distortable Canvas, Abstracted Multi-Level Gradient Descent (AMGD), Data-Efficient Innovation*

1. INTRODUCTION

Artificial Intelligence (AI) has revolutionized science, industry, and daily life, driving advancements in automation, creativity, and decision-making [1]. However, most AI systems still rely on vast amounts of annotated data, which limits their scalability and applicability in data-scarce environments [2]. This dependency has motivated researchers to seek more data-efficient learning paradigms that better emulate how humans learn from limited experience. Unlike machines, humans can acquire new concepts from only one or a few examples, generalize them flexibly, and apply abstract knowledge to novel situations [3][4].

Recent studies have explored the cognitive foundations of such abilities, showing that human learning involves hierarchical reasoning, analogical mapping, and structural abstraction rather than statistical memorization [5]. Inspired by these mechanisms, one-shot learning has emerged as a

promising approach that aims to replicate human-like adaptability within artificial systems. Yet, despite significant progress, existing models often remain opaque and domain-restricted, making it difficult to interpret or transfer their learned representations across tasks.

To overcome these challenges, this paper proposes an Interdisciplinary One-Shot Learning Framework that mimics human cognition to achieve data efficiency, interpretability, and cross-domain generalization. Grounded in human-inspired perception and abstraction, the framework introduces a transparent mechanism for similarity reasoning and adaptive optimization. By integrating perspectives from design, engineering, business, and culture, this work advances the vision of human-like artificial intelligence—one that learns efficiently, reasons transparently, and innovates beyond disciplinary boundaries [6].

2. RELATED WORK

Recent advances in machine learning have driven remarkable progress in perception, generation, and reasoning; however, these achievements often depend on massive labeled datasets, making them impractical in many real-world scenarios [7]. To address this challenge, researchers have explored a range of data-efficient learning paradigms, including One-Shot Learning, Few-Shot Learning, Meta-Learning, and Transfer Learning [8].

One-Shot Learning aims to enable models to learn new concepts from only one or a few examples, a capability that closely resembles human cognition [9]. Classic approaches, such as Siamese and Matching Networks, introduced metric-based architectures that measure similarity between instances rather than relying on large-scale classification [10]. This paradigm highlights the importance of feature generalization and cross-class similarity, paving the way for more human-like learning behavior.

Few-Shot Learning extends the one-shot paradigm by training models to adapt quickly with only a small number of examples [11]. It focuses on learning robust priors or meta-knowledge that can generalize to unseen classes. Notable works such as Prototypical Networks demonstrate how

embedding-based similarity spaces can achieve competitive performance under extreme data scarcity [12].

Meta-Learning, or “Learning to Learn,” provides a broader framework for data-efficient adaptation [7]. Instead of learning specific tasks, meta-learning systems learn how to optimize themselves for new tasks based on prior experiences [13]. The Model-Agnostic Meta-Learning (MAML) algorithm, for example, learns initial model parameters that can be rapidly fine-tuned to new problems [14]. Such mechanisms make meta-learning foundational to achieving human-like flexibility in artificial systems [15].

Transfer Learning further addresses data scarcity by reusing knowledge acquired from one domain to improve performance in another [16]. Foundational work in this field explores how feature representations learned from large datasets can be fine-tuned for related but smaller tasks. Theoretical studies have also investigated conditions under which transfer leads to positive or negative effects. In interdisciplinary contexts, transfer learning enables efficient cross-domain adaptation, facilitating innovation across diverse fields such as natural language processing, computer vision, and engineering.

Despite these advances, existing data-efficient paradigms often face limitations in interpretability and cross-domain generalization. Deep neural models can achieve high accuracy but frequently operate as “black boxes,” making their decision processes difficult to explain [16]. This opacity limits their reliability in domains that demand transparency, such as healthcare, business strategy, or cultural analysis. Addressing these issues requires frameworks that not only learn efficiently from limited data but also provide cognitively meaningful explanations of how knowledge is transferred and represented.

Human intelligence is inherently data-efficient and generalizable, enabling people to form abstractions and analogies from minimal experience. Research in human-like learning aims to model this ability computationally, bridging cognitive science and machine learning. Theoretical and empirical studies show that humans rely on hierarchical abstraction and analogy-making to construct conceptual knowledge [17]. Modeling these mechanisms has become central to the pursuit of human-level artificial intelligence.

Interdisciplinary applications of AI further test the limits of generalization. Many real-world challenges—such as product design, systems engineering, market forecasting, and cultural heritage analysis—require integrating insights across heterogeneous data and conceptual domains [18]. Traditional AI models are typically domain-specific, which constrains their ability to transfer knowledge across disciplines. Recent research has demonstrated how transfer learning and meta-learning can facilitate cross-domain innovation, for instance in adaptive design systems and intelligent business analytics [19]. These studies reveal that when models capture shared structural representations, they can bridge conceptual gaps between domains and promote creativity.

However, major gaps remain in building interpretable and transparent interdisciplinary systems. Many human-like learning models still fall short in replicating the flexibility and explainability of human cognition. Similarly, interdisciplinary AI applications often require manual adaptation to each field, preventing scalable generalization [20]. To overcome these barriers, the present study introduces a cognitively inspired Interdisciplinary One-Shot

Learning Framework that combines visual-intuitive similarity modeling and multi-level optimization, thereby enhancing both interpretability and data efficiency [21].

3. METHODOLOGY AND SYSTEM DESIGN

Our framework is designed to achieve human-like, data-efficient learning by combining cognitively inspired similarity modeling with adaptive optimization. It integrates two core components—the Distortable Canvas and the Abstracted Multi-Level Gradient Descent (AMGD)—to emulate human perception of general appearance similarity and hierarchical reasoning [22].

3.1. *The Distortable Canvas: Modeling General-Appearance Similarity*

Humans perceive similarity not by comparing isolated features, but by assessing the overall structure and how one object could transform into another with minimal effort. The Distortable Canvas formalizes this principle by representing images or visual patterns as deformations on a flexible surface [23]. The similarity between two instances is defined by the minimal “energy” required to transform one into another, considering both geometric deformation and photometric variation.

The distortable canvas operates by finding the optimal transformation that maps one image onto another, minimizing a combined distortion cost. This cost function considers both the geometric deformation of the canvas (D_V) and the photometric changes in

color or intensity (D_C) [24]. The objective is to find a transformation function, T , that minimizes:

$$\text{Cost}(I_1, I_2) = \min_T [\alpha * DV(T) + \beta * DC(I_1, T(I_1), I_2)] \quad (1)$$

where I_1 and I_2 are the two images being compared, α and β are weighting parameters,

$DV(T)$ quantifies the energy required for the geometric transformation T , and $DC(I_1, T(I_1), I_2)$ measures the color difference between the transformed I_1 and I_2 [25].

The transformation T is not restricted to affine or projective transformations but can

encompass highly non-linear deformations, reflecting the complex ways humans mentally manipulate visual information. The output of this process is a similarity score, or distance, that quantifies how ‘close’ two images are in this cognitively-inspired space [26] [27].

3.2. *Abstracted Multi-level Gradient Descent (AMGD) for Optimization*

Optimizing transformations in high-dimensional visual spaces is challenging due to numerous local minima. To overcome this, we introduce AMGD, a hierarchical optimization algorithm inspired by human problem-solving strategies. AMGD initiates with coarse-level transformations to capture global structure and progressively refines them at finer levels, adjusting learning rates and update strategies according to the complexity encountered at each stage [28]. This multi-level process allows the algorithm to avoid local traps, converge more efficiently, and mirror human reasoning, where understanding typically begins with the overall concept before focusing on detailed features. By combining coarse-to-fine updates with adaptive adjustments, AMGD produces transformations that are both accurate and

interpretable, revealing the trajectory of perceptual alignment between instances.

3.3. Component Details

The framework is structured as a modular system comprising the Distortable Canvas, the AMGD optimizer, and a domain-specific application layer. The learned similarity space serves as a flexible foundation, enabling its application across diverse domains. In design, the system facilitates the generation of sketches aligned with reference images, models user preferences from limited interactions, and supports adaptive design systems capable of responding to evolving trends [29]. In engineering, it accelerates prototyping from sparse data, predicts potential system failures through anomaly detection, and expedites product development [30]. For business applications, the framework can identify emerging market trends from noisy or limited datasets, model consumer behavior for personalized recommendations, and support data-driven product innovation. In cultural and heritage domains, it enables automated recognition of artistic styles, paleographic analysis of historical manuscripts, and the preservation of intangible cultural assets by capturing essential patterns from limited examples. The implementation leverages Python for core logic with libraries such as NumPy and SciPy for numerical operations, while PyTorch or TensorFlow are used for efficient image handling and GPU acceleration. Custom data structures manage transformation parameters and the AMGD hierarchy, and the modular design ensures scalability to larger datasets or more complex transformations. This methodology demonstrates a robust, interpretable, and cognitively grounded framework capable of addressing diverse real-world challenges through data-efficient one-shot learning.

4. EXPERIMENTS AND RESULTS

To rigorously evaluate the efficacy and interdisciplinary applicability of our proposed one-shot learning framework, we conducted a series of experiments across various benchmarks and real-world scenarios relevant to design, engineering, business, and cultural applications. Our experimental design emphasizes the framework's ability to learn from extremely limited data, its interpretability, and its performance against conventional methods in data-scarce environments. All experiments were conducted without any pre-training or data augmentation, adhering strictly to the 'only-few-shot' paradigm.

4.1. Experimental Setup and Benchmarks

We utilized a combination of established datasets and synthetically constructed data tailored to fit specific interdisciplinary challenges. For evaluating the core learning mechanism (Distortable Canvas and AMGD), we selected primary benchmarks from the MNIST, EMNIST, Omniglot, and QuickDraw datasets—drawing on the evaluation

protocols of prior work in the field while refining task parameters to focus specifically on the one-shot and few-shot regimes. For interdisciplinary applications, we designed specific experimental setups:

Design Application (Sketch Recognition & Style Transfer): We used a subset of the QuickDraw dataset (for sketch recognition) and a curated set of design mood

boards and product images (for style transfer, represented as image sets). The task involved recognizing unseen sketches after exposure to only one example per category and transferring stylistic elements from a single reference image.

Engineering Application (Material Microstructure Classification): We generated synthetic images representing various material microstructures (e.g., different grain sizes, phase distributions) under controlled conditions. The challenge was to classify novel microstructures based on a single training image per class, mimicking rapid material characterization in R&D.

Business Application (Market Trend Anomaly Detection): We constructed time-series data representing simplified market trends (e.g., stock price fluctuations, sales volumes) with embedded subtle anomalies. The task was to detect these anomalies after learning 'normal' patterns from a few historical data points.

Performance metrics included accuracy, F1-score (for classification tasks), and a qualitative assessment of interpretability (e.g., visual analysis of transformation flows). For all classification tasks, we report mean and standard deviation from 5 independent runs to ensure statistical robustness. Comparison models included traditional nearest-neighbor classifiers, Support Vector Machines (SVMs), and simplified neural network architectures trained from scratch on the same limited data, without pre-training.

4.2. Core Learning Performance on Abstract Visual Tasks

Our framework's core ability to learn general-appearance similarity from minimal data was first validated on abstract visual tasks, replicating and extending the findings from prior work. The results consistently demonstrate superior performance in the 'only-few-shot' regime.

For the MNIST dataset (handwritten digits), we evaluated the model's accuracy when trained with only N examples per class ($N=1, 2, 3, 4$). As shown in Figure 1, our model significantly outperforms both traditional ML algorithms and contemporary neural networks in this extreme data scarcity setting. With just one training image per class, our model achieved an average accuracy of 80.2% ($\pm 1.5\%$), which increased to 90.5% ($\pm 0.8\%$) with four training images per class. This highlights the framework's exceptional data efficiency and generalization capabilities.




Figure 1. MNIST Only-Few-Shot Classification Accuracy

The Omniglot dataset, known for its diverse handwritten characters from 50 different alphabets, presents a more challenging one-shot learning task. Our model achieved a remarkable 93.2% ($\pm 0.7\%$) accuracy in the 20-way one-shot

classification task, approaching human-level performance (95.5%). This result, depicted in Figure 2, is particularly significant as our model did not utilize any background set for pre-training or stroke information, unlike many state-of-the-art methods designed specifically for Omniglot.




Figure 2. Omniglot One-Shot Classification Performance (Error Rate Leaderboard)

4.3. Interdisciplinary Application Results

Beyond abstract visual tasks, we demonstrated the framework's versatility and effectiveness in solving real-world problems across the four interdisciplinary domains.

Our framework proved highly effective in recognizing novel sketches and facilitating style transfer with minimal examples. For sketch recognition on a subset of QuickDraw, the model achieved 88.7% ($\pm 1.2\%$) accuracy with only one training example per sketch category, demonstrating its ability to capture the essence of a drawing from a single instance. This is crucial for adaptive design systems where new design elements are constantly introduced.

For style transfer, we qualitatively assessed the generated outputs. Given a single reference image (e.g., a mood board or a product with a specific aesthetic) and a target image (e.g., a rough design sketch), our model successfully applied the stylistic elements of the reference to the target, producing visually coherent and aesthetically pleasing results.

In the engineering domain, our framework demonstrated robust performance in classifying material microstructures from limited data. Using synthetically generated images of various material phases and grain structures, the model achieved 91.3% ($\pm 0.9\%$) accuracy in classifying unseen microstructures with only one training image per class. This capability is vital for rapid material characterization and quality control in manufacturing processes where obtaining

numerous samples for each new material variant is impractical.

Figure 3 presents the confusion matrix for the material microstructure classification, highlighting the model's ability

to differentiate between subtle variations in material composition and structure. The high diagonal values indicate strong classification performance across all classes.





Figure 3. Confusion Matrix for One-Shot Material Microstructure Classification

Furthermore, the interpretability of our model allowed engineers to visualize the 'transformation flow' between different microstructures, providing insights into the key structural differences that the model learned. This visual explanation, shown in Figure 4, can aid in understanding material properties and optimizing manufacturing parameters.

Figure 4. Visualization of Transformation Flow Between Different Material Microstructures

For business applications, our framework excelled in detecting subtle anomalies in market trend data. By representing time-series data as 'visual patterns' (e.g., using Gramian Angular Fields or Recurrence Plots), the model learned normal market behavior from a few historical examples. When presented with new data, it accurately identified anomalous patterns indicative of market shifts or unusual events.

Figure 5 illustrates the anomaly detection performance, showing the reconstruction error for normal versus anomalous market patterns. Anomalous patterns consistently exhibited significantly higher reconstruction errors, enabling effective detection. The model achieved an F1-score of 0.89 (± 0.03) for anomaly detection, demonstrating its reliability in identifying critical business events from sparse data.






Figure 5. Reconstruction Error for Normal vs. Anomalous Market Trend Patterns

In the cultural domain, our framework proved highly effective in recognizing historical script characters, a challenging task due to the uniqueness and rarity of many historical scripts. On a curated dataset of historical script characters, the model achieved 85.1% ($\pm 1.1\%$) accuracy in recognizing unseen characters with only one training example per character class. This capability is invaluable for digital humanities and cultural heritage preservation efforts.

4.4. Interpretability and Generalization Analysis

A key advantage of our framework is its inherent interpretability, stemming from the white-box nature of the

Distortable Canvas and the AMGD optimizer. Unlike opaque deep learning models, our framework allows for direct visualization of the learned similarity space and the transformation flows between data points. This transparency provides profound insights into how the model makes decisions and generalizes from limited examples.

Figure 6 presents a t-SNE visualization of the learned similarity space for a multi-class dataset, demonstrating how our framework effectively clusters similar instances while maintaining clear separation between different classes, even with minimal training data. The clusters are formed based on general-appearance similarity, reflecting human intuition.




Figure 6. t-SNE Visualization of the Learned Similarity Space

Furthermore, the transformation flows provide a direct visual explanation of the model's reasoning. When the model

identifies two instances as similar, the transformation flow illustrates the minimal 'distortion' required to morph one into

the other. This not only validates the model's decisions but also offers new perspectives on visual relationships, akin to human insights. This level of interpretability is critical for high-stakes applications and for fostering trust in AI systems.

4.5. Comparison with State-of-the-Art Methods

While direct comparisons with all state-of-the-art (SOTA) few-shot learning methods are challenging due to their reliance on extensive pre-training or specific meta-training datasets, we conducted comparative analyses against representative methods in the 'from-scratch' or 'minimal pre-

training' categories. Our framework consistently demonstrated superior or competitive performance, particularly in scenarios with extreme data scarcity and without any external pre-trained knowledge.

Table 1 summarizes the performance comparison across various benchmarks, highlighting our model's advantages in data efficiency and interpretability. Our model's ability to achieve high accuracy with only one or a few examples, without the need for massive pre-training, positions it as a highly promising alternative for real-world applications where data is a limiting factor.

TABLE I. PERFORMANCE COMPARISON WITH REPRESENTATIVE FEW-SHOT LEARNING METHODS

Method	MNIST (1-shot)	Omniglot (1-shot, 20-way)	Material Microstructure (1-shot)	Market Anomaly (F1-score)	Interpretability	Pre-training Required
Our Framework	80.2%	93.2%	91.3%	0.89	High (White-box)	No
Prototypical Networks (from scratch)	65.1%	88.5%	72.8%	0.75	Medium	No
MAML(from scratch)	68.9%	89.1%	75.2%	0.78	Medium	No
Simple k- NN	45.0%	60.0%	55.0%	0.60	High	No
TextCaps (with pre training)	N/A	N/A	N/A	N/A	Low	Yes

These results collectively underscore the robustness, versatility, and interpretability of our interdisciplinary one-shot learning framework. By mimicking human cognitive processes, we have developed a powerful tool that can effectively address data scarcity challenges and foster innovation across a wide array of scientific and practical domains.

5. ANALYSIS AND DISCUSSION

The experimental results presented in Section 4 provide compelling evidence for the efficacy, interpretability, and interdisciplinary applicability of our one-shot learning framework. This section delves deeper into the implications of these findings, discusses the observed phenomena, compares our approach with related work, highlights the research value, and acknowledges the limitations and potential sources of error.

5.1. Interpretation of Results and Observed Phenomena

The results demonstrate that our framework achieves remarkable performance under extremely limited data conditions across both abstract visual tasks, such as MNIST and Omniglot, and a range of applied domains including design, engineering, business, and cultural studies. This efficiency can be attributed to the cognitively inspired modeling of similarity through the Distortable Canvas, which captures intrinsic relationships between instances by minimizing perceptual distortions rather than relying on conventional feature-based metrics, allowing robust generalization from very few examples. Coupled with the hierarchical optimization strategy of Abstracted Multi-Level Gradient Descent, which progressively abstracts and refines transformations, the framework effectively navigates complex similarity spaces, avoiding local minima and converging on stable solutions in a manner analogous to human learning, where global understanding precedes detailed refinement. Importantly, the framework provides transparent insight into its decision-making process through visualization of transformation flows, offering interpretability that is particularly valuable in high-stakes contexts, such as material microstructure analysis, medical

diagnostics, and engineering systems, where understanding the rationale behind outcomes is as critical as the results themselves. These properties collectively illustrate that effective learning can emerge from cognitively grounded similarity perception rather than large amounts of data.

5.2. Comparison with Related Work

Compared with existing few-shot, meta-learning, and transfer learning methods, the framework distinguishes itself by achieving high performance without extensive pre-training or meta-datasets, highlighting its genuine data efficiency. Furthermore, its emphasis on human-inspired general-appearance similarity and white-box interpretability provides actionable insight that many black-box approaches lack, fostering trust and facilitating understanding across domains. Beyond empirical performance, this approach advances the theoretical understanding of human-like learning, demonstrating that generalization can arise from robust similarity perception. Practically, it offers a versatile tool for innovation in data-scarce contexts, enabling rapid design prototyping, predictive maintenance, market analysis, and the preservation of cultural heritage. Nonetheless, challenges remain, including computational costs for high-dimensional transformations, optimal representation of non-visual data, scalability to very large datasets, and generalization to entirely novel tasks. Developing quantitative metrics for interpretability also represents an important direction for future research. Overall, the framework provides a foundation for more human-like, interpretable, and data-efficient learning, offering both theoretical insight and practical utility across disciplines.

6. CONCLUSION

This paper presents a novel interdisciplinary one-shot learning framework that rethinks how learning and generalization can occur under extreme data scarcity, drawing inspiration from human cognitive processes. By conceptualizing a distortable canvas to model general-appearance similarity and optimizing this space through an Abstracted Multi-level Gradient Descent (AMGD) algorithm, we developed a transparent, white-box model that combines

remarkable data efficiency with inherent interpretability. Unlike conventional machine learning paradigms, which typically rely on large datasets and often opaque decision processes, our approach enables high-performance learning in an ‘only-few-shot’ regime without requiring pre-training or data augmentation, providing a robust solution in contexts where data is severely limited.

The interpretability of the framework is a particularly notable strength, as the visualization of transformation flows and the learned similarity space offers direct insight into the reasoning process, allowing users to understand how the model arrives at its conclusions. This transparency not only fosters trust but also facilitates error analysis and the extraction of meaningful knowledge from the learned representations, which is especially valuable in high-stakes applications such as material design, engineering diagnostics, and cultural heritage preservation. The framework’s ability to generalize across diverse interdisciplinary domains demonstrates its versatility, highlighting its potential as a unifying mechanism for innovation that bridges traditionally disparate fields. By supporting applications ranging from adaptive design and agile engineering to strategic business intelligence and preservation of cultural artifacts, the framework illustrates how cognitively inspired approaches can enable both practical utility and methodological insight.

Moreover, by emulating human processes of generalization and abstraction, the framework contributes to the broader goal of developing more human-like learning systems. It demonstrates that effective learning and robust generalization can emerge from cognitively grounded similarity perception rather than relying solely on large-scale statistical accumulation, thereby advancing our understanding of fundamental principles underlying human cognition. At the same time, the framework provides a foundation for future exploration in several directions, including extending its applicability to more complex and heterogeneous data types, enhancing computational efficiency, and developing quantitative metrics for interpretability. In doing so, it establishes a pathway for creating learning systems that are not only data-efficient and interpretable but also capable of adapting to novel problems in dynamic and resource-constrained environments.

REFERENCES

- [1] Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33, 1877-1901.
- [2] Bansal, M. A., Sharma, D. R., & Kathuria, D. M. (2022). A systematic review on data scarcity problem in deep learning: solution and applications. *ACM Computing Surveys (Csur)*, 54(10s), 1-29. <https://doi.org/10.1145/3502287>
- [3] Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature machine intelligence*, 1(5), 206-215. <https://doi.org/10.1038/s42256-019-0048-x>
- [4] Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *science*, 331(6022), 1279-1285. <https://doi.org/10.1126/science.1192788>
- [5] Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, 350(6266), 1332-1338. <https://doi.org/10.1126/science.aab3050>
- [6] Vinyals, O., Blundell, C., Lillicrap, T., & Wierstra, D. (2016). Matching networks for one shot learning. *Advances in neural information processing systems*, 29.
- [7] Wirth, K. R., & Perkins, D. (2008, September). Learning to learn.
- [8] Snell, J., Swersky, K., & Zemel, R. (2017). Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30.
- [9] Koch, G., Zemel, R., & Salakhutdinov, R. (2015, July). Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop (Vol. 2, No. 1, pp. 1-30)*.
- [10] Finn, C., Abbeel, P., & Levine, S. (2017, July). Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning (pp. 1126-1135)*. PMLR.
- [11] Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction (Vol. 1, No. 1, pp. 9-11)*. Cambridge: MIT press.
- [12] Thota, S., Chitta, S., Vangoor, V. K. R., Ravi, C. S., & Bonam, V. S. M. (2023). *Few-Shot Learning in Computer Vision: Practical Applications and Techniques*. Human-Computer Interaction, 3(1).
- [13] Yosinski, J., Clune, J., Bengio, Y., & Lipson, H. (2014). How transferable are features in deep neural networks?. *Advances in neural information processing systems*, 27.
- [14] Alyafeai, Z., AlShaibani, M. S., & Ahmad, I. (2020). A survey on transfer learning in natural language processing. *arXiv preprint arXiv:2007.04239*. <https://doi.org/10.48550/arXiv.2007.04239>
- [15] Pan, S. J., & Yang, Q. (2009). A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10), 1345-1359. <https://doi.org/10.1109/TKDE.2009.191>
- [16] Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., & Vaughan, J. W. (2010). A theory of learning from different domains. *Machine learning*, 79(1), 151-175. <https://doi.org/10.1007/s10994-009-5152-4>
- [17] Zhang, W., Deng, L., Zhang, L., & Wu, D. (2022). A survey on negative transfer. *IEEE/CAA Journal of Automatica Sinica*, 10(2), 305-329. <https://doi.org/10.1109/JAS.2022.1060004>
- [18] Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and brain sciences*, 40, e253. <https://doi.org/10.1017/S0140525X16001837>
- [19] Sloutsky, V. M. (2003). The role of similarity in the development of categorization. *Trends in cognitive sciences*, 7(6), 246-251. [https://doi.org/10.1016/S1364-6613\(03\)00109-8](https://doi.org/10.1016/S1364-6613(03)00109-8)
- [20] Pinker, S. (2004). *The blank slate: The modern denial of human nature*. New York.
- [21] Yu, H., Mineyev, I., Varshney, L. R., & Evans, J. A. (2025). Learning from one and only one shot. *npj Artificial Intelligence*, 1(1), 13. <https://doi.org/10.1038/s44387-025-00017-7>
- [22] Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and brain sciences*, 40, e253. <https://doi.org/10.1017/S0140525X16001837>
- [23] Hassabis, D., Kumaran, D., Summerfield, C., & Botvinick, M. (2017). Neuroscience-inspired artificial intelligence. *Neuron*, 95(2), 245-258. <https://doi.org/10.1016/j.neuron.2017.06.011>
- [24] Pena, M. L. C., Carballal, A., Rodríguez-Fernández, N., Santos, I., & Romero, J. (2021). Artificial intelligence applied to conceptual design. A review of its use in architecture. *Automation in Construction*, 124, 103550. <https://doi.org/10.1016/j.autcon.2021.103550>
- [25] Enholm, I. M., Papagiannidis, E., Mikalef, P., & Krogstie, J. (2022). Artificial intelligence and business value: A literature review. *Information systems frontiers*, 24(5), 1709-1734. <https://doi.org/10.1007/s10796-021-10186-w>
- [26] De Souza, V. L. T., Marques, B. A. D., Batagelo, H. C., & Gois, J. P. (2023). A review on generative adversarial networks for image generation. *Computers & Graphics*, 114, 13-25. <https://doi.org/10.1016/j.cag.2023.05.010>
- [27] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *nature*, 521(7553), 436-444. <https://doi.org/10.1038/nature14539>
- [28] Zhang, W., Gu, X., Tang, L., Yin, Y., Liu, D., & Zhang, Y. (2022). Application of machine learning, deep learning and optimization algorithms in geoenvironment and geoscience: Comprehensive review and future challenge. *Gondwana Research*, 109, 1-17. <https://doi.org/10.1016/j.gr.2022.03.015>
- [29] Chollet, F. (2017). Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer*

vision and pattern recognition (pp. 1251-1258).
<https://doi.org/10.48550/arXiv.1610.02357>

[30] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 770-778).