Engineering and Technology Journal e-ISSN: 2456-3358

Volume 10 Issue 03 March-2025, Page No.- 4224-4230

DOI: 10.47191/etj/v10i03.31, I.F. – 8.482

From Shortcut Solutions to First Principles: Addressing Challenges and Cultivating Innovation in AI Research

Ayse Kok Arslan

Researcher, Silcion Valley Chapter-Oxford ALumni

ABSTRACT: The rapid progress in artificial intelligence (AI) has led to an increased use of shortcut methods like distillation-based training and prompt optimization. While these approaches offer quick performance improvements, they may hinder long-term innovation in the field.

This study highlights the importance of a balanced research approach that combines short-term performance goals with long-term advancements in search and inference optimization. It also suggests redesigning educational systems to promote deep understanding, curiosity, and critical thinking. The paper concludes that while developing intelligent AI systems is important, nurturing a new generation of researchers who can think from first principles is crucial for sustainable AI advancement and innovation.

INTRODUCTION

AI research has experienced a dramatic shift with the emergence of models capable of complex reasoning. However, the focus on performance gains over transparent innovation has rendered traditional research methods inadequate for addressing the complexities of modern AI projects.

The paper examines the challenges of shortcut approaches in AI, considering both their technical limitations and broader impacts on research culture and future AI researchers. A key concern is the decline in first-principles thinking, which is crucial for sustainable AI innovation. The discussion covers:

1. Main paradigms in large language model (LLM) training, using the "shortcut vs. journey learning" framework by Qin et al. (2024)
2. Primary reasoning methods in LLMs
3. Overview of the Technical Transparency Index for evaluating LLM capabilities
4. Lessons learned from current practices

The paper concludes with recommendations for a balanced research portfolio that combines short-term performance goals with long-term advancements in search and inference optimization.

Main Paradigms in Machine Learning

The recent rise of artificial intelligence (AI) technologies has brought with it a surge in methodologies focused on rapid performance enhancements. On the other hand, several short-cut methods such as distillation-based training and prompt optimization, emerged to deliver immediate results without any concerns about their long-term implications. These "shortcut methods," prioritize efficiency at the expense of foundational advancements, potentially leading to a performance ceiling effect and missed opportunities for transformative innovation.

Qin et al. (2024) highlight a prevailing issue in contemporary machine learning and large language model training, which they refer to as "shortcut learning." While shortcut learning has driven significant progress, its inherent limitations underscore its inability to produce truly robust and adaptable AI systems.

As seen in Figure 1, this approach, commonly associated with supervised fine-tuning, is characterized by several critical limitations:

1. **Quick Results Orientation**: Emphasis is placed on achieving specific performance metrics or completing narrowly defined tasks within minimal timeframes.
2. **Heavy Data Dependency**: Performance gains are frequently reliant on expanding the volume of training data, rather than improving the underlying algorithms.
3. **Limited Generalization**: Models exhibit sharp performance declines when applied to scenarios beyond the distribution of their training data.
4. **Lack of Self-Correction**: Current systems often fail to identify and rectify their own errors, rendering them inadequate for addressing the complexities of real-world challenges.

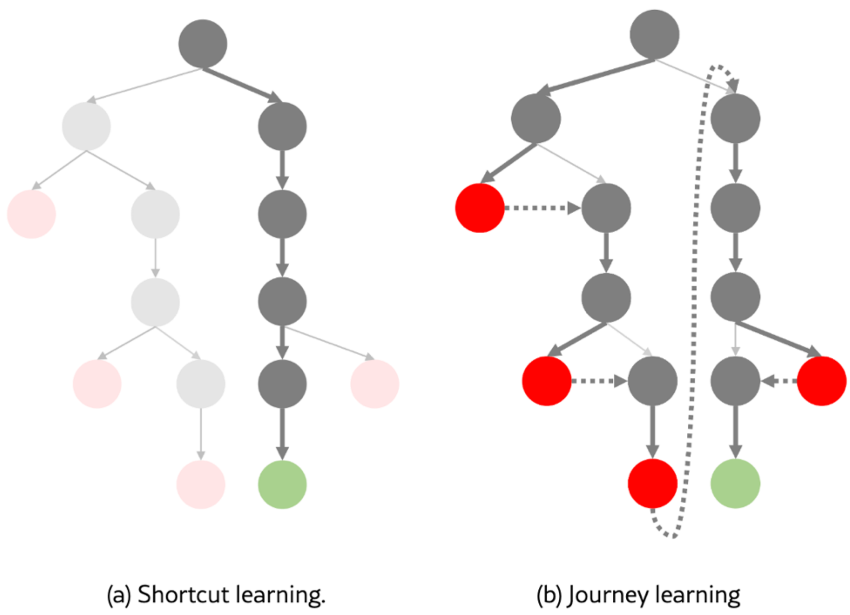

Fig.1 Main Paradigms in AI Training

Qin et al. (2024) argue that the implications of this approach extend beyond technical inefficiencies to systemic challenges within AI research and education as seen in Figure 2:

1. **Transparency Deficits**: The lack of transparency in methodology reporting complicates efforts to validate and build upon claimed advancements, often distorting the broader understanding of progress within the field.
2. **Innovation Stagnation**: Increasing reliance on existing powerful models shifts focus from developing fundamentally novel techniques to refining prompt engineering, which may hinder long-term advancements.
3. **Model Limitations**: Distillation-based training confines models to the capabilities of their predecessors, creating a "ceiling effect" that impedes breakthroughs in new domains.
4. **Educational Gaps**: The current paradigm neglects opportunities to nurture foundational research and problem-solving skills among the next generation of AI researchers.

Moreover, Qin et al. (2024) advocate for a comprehensive reimagining of the research ecosystem in the AI era, emphasizing the following pillars:

1. **Addressing Challenges in Modern AI Research**: The rapid evolution of AI necessitates a shift towards greater transparency and real-time feedback within long-term, team-based research endeavors, thereby maintaining researcher motivation and enhancing information dissemination.
2. **Fostering Open Science**: Encouraging the sharing of not just trained models but also the datasets, tools, and methodologies underpinning their development fosters collective progress and replicability.
3. **Laying Foundations for AI in Scientific Discovery**: Detailed documentation of research processes, including failures, creates invaluable datasets for training AI models capable of understanding and replicating scientific methodologies.
4. **Promoting Responsible AI Development**: A commitment to transparency and ethical research practices establishes public trust and sets a high standard for responsible AI innovation.

By adopting these principles, the field of AI can transition from a narrow focus on performance metrics to a more holistic and sustainable paradigm of innovation.

Acknowledging the limitations of current methodologies, Qin et al. (2024) propose a transformative framework known as "journey learning." Journey learning is designed to facilitate continuous progression in AI systems through iterative cycles of learning, reflection, backtracking, and adaptation. By fostering these capabilities, it seeks to cultivate higher levels of intelligence and flexibility in AI systems.

Journey learning transcends conventional learning methods, introducing a paradigm that mirrors human-like cognitive processes. It promises the creation of more versatile, human-like AI systems, capable of meaningful interaction and application across diverse domains.

Characteristic	Shortcut Learning	Journey Learning
Learning Depth	Surface features and simple correlations	Deep causal relationships and underlying principles
Reasoning Ability	Limited, struggles with complex reasoning	Powerful, demonstrates human-like reasoning
Self-Improvement	Lacks self-correction mechanisms	Continuous self-assessment and improvement
Generalization	Limited, easily affected by data distribution changes	Strong, can handle new situations
Innovation Capacity	Limited, struggles to solve new problems	High, can generate innovative solutions
Data Dependency	Highly dependent on large training datasets	More focused on quality and learning strategies
Interpretability	Poor, often seen as a "black box"	Better, can track internal reasoning processes
Ethical Considerations	May unintentionally amplify data biases	Easier to implement ethical constraints and adjustments
Security	Vulnerable to adversarial attacks	More robust, able to identify potential threats
Long-term Value	Quick results in specific tasks	Paves the way for AGI development
Human Analogy	Exam-oriented education, crash courses	Comprehensive education, lifelong learning

Table 2: Comparison between Shortcut Learning and Journey Learning.

Unlike shortcut learning, this paradigm envisions the development of AI systems that go beyond narrow, task-specific applications to become adaptable, reasoning entities capable of addressing real-world challenges with sophistication and nuance.

Building on these innovative approaches to learning, Qin et al. (2024) outline a broader framework for scientific exploration and collaboration in the AI era. Their framework emphasizes four foundational aspects:

1. **Addressing Challenges in Modern AI Research**: Due to the rapid advancement of AI technologies, research projects often involve extended, collaborative endeavors which present unique challenges, including restricted information flow within the scientific community and diminished researcher motivation due to delayed gratification. To counter these issues, the proposed framework advocates for enhanced transparency, real-time feedback mechanisms, and recognition systems that sustain commitment to long-term research efforts.
2. **Fostering Open Science and Collective Advancement**:
   Derived from the principles of open science, the framework emphasizes the importance of sharing both trained AI models and the documentation of the tools, datasets, and methodologies employed. This approach fosters a collective progress, enabling researchers to build upon existing knowledge and innovations effectively.
3. **Laying the Foundation for AI-Driven Scientific Discovery**:
   Meticulous documentation of the scientific exploration process, including both successes and failures, is pivotal. Such records create a valuable dataset for training AI models to understand and replicate scientific methodologies. By capturing the entirety of the research journey, this framework aligns AI development with foundational scientific principles, enabling more robust and insightful discoveries.
4. **Promoting Responsible AI Development**: Transparent documentation of research processes and decision-making sets a high standard for accountability and ethical AI practices. This transparency is critical for building public trust and a research culture that prioritizes responsible innovation. By adopting a holistic approach, the framework encourages long-term sustainability in AI research and development.

In conclusion, journey learning represents a significant evolution in AI methodology, addressing the shortcomings of shortcut learning while establishing a foundation for collaborative, responsible, and impactful scientific exploration in the AI era. This paradigm holds significant potential to redefine the boundaries of what AI systems can achieve and how they integrate into society.

Review of Existing Work on LLM Reasoning Methods

Foundational models employ various techniques to construct long chains of reasoning necessary for solving complex problems. These chains often integrate reflection, error correction, and backtracking steps, with each method offering trade-offs between computational efficiency and the thoroughness of reasoning.

Method I: Complete Human Thought Process Annotation

Human problem-solving is rarely linear, often involving reflection, backtracking, and iterative revisions in the face of challenges. This natural problem-solving process mirrors the characteristics of long-chain reasoning. By meticulously documenting human strategies, researchers can generate authentic, comprehensive training data to enhance AI reasoning capabilities.

Method II: Multi-Agent Approach

In contrast to journey learning, where policy models do not directly respond to feedback, a multi-agent framework involves assigning distinct roles to different agents. For instance, a multi-agent debate system can feature a policy model generating reasoning chains while a critique model evaluates whether to proceed or backtrack. This dynamic interaction naturally produces high-quality training data as solutions are collaboratively refined.

Method III: Distillation from Advanced Models

Advanced AI models, known for their robust reflection and self-correction abilities, serve as valuable mentors in the distillation process. Weaker models learn from the outputs of stronger models through careful prompting. However, access to internal reasoning processes is often restricted, necessitating meticulous design to ensure effective knowledge transfer.

To further advance reasoning capabilities, several innovative methods have been introduced:

1. **Process-Level Reward Models (PRMs)**
   Process reward models evaluate responses from large language models (LLMs) at a granular level, particularly in mathematical reasoning (Lightman et al., 2024; Uesato et al., 2022; Xia et al., 2024). Techniques such as Monte Carlo Tree Search model multi-step reasoning as a Markov Decision Process (Silver et al., 2016), enabling online (Chen et al., 2024) and offline (Wang et al., 2024c) reasoning improvements.
2. **Chain-of-Thought (CoT) Theory**
   CoT prompting enhances reasoning by incorporating intermediate reasoning steps, significantly improving performance on tasks like arithmetic and commonsense reasoning (Wei et al., 2022). Recent work integrates error-correction data during pretraining, yielding higher accuracy without the need for multi-round prompting (Ye et al., 2024).
3. **Internal Thought**
   Researchers have emphasized reflective and iterative processes within AI models. Early methods like STaR (Zelikman et al., 2022) enabled models to generate rationales that iteratively refine their outputs. Extensions like Quiet-STaR (Zelikman et al., 2024a) train models to predict and explain text step-by-step, fostering deeper reasoning capabilities.
4. **Inference Time Scaling**
   Scaling inference time offers an alternative to traditional methods of expanding model parameters or data volume (Sardana and Frankle, 2023; Snell et al., 2024). Benefits include resource efficiency, adaptable computation for complex tasks, and enhanced reasoning through iterative problem-solving.
5. **Search-to-Thought**
   Modern approaches have shifted from explicit search-based methods, such as alpha-beta pruning (Campbell et al., 2002), to implicit reasoning using internal model states. Deng et al. (2023) demonstrate how distilling intermediate steps from teacher models enables efficient task-solving, reducing reliance on computationally expensive algorithms.

To mitigate risks like model collapse, researchers suggest balancing human-authored and LLM-generated data during training (Gerstgrasser et al., 2024). In order to refine LLM reasoning and reduce hallucination risks, the following strategies are recommended:

1. **Enhanced Data Granularity**: Problems are broken into finer, digestible steps to ensure comprehensive understanding at each stage.
2. **Gradual Reasoning**: Models frequently pause to reflect, mirroring human thinking patterns, thereby improving engagement with the reasoning process.
3. **Student-Explorer Perspective**: Adopting a tone of discovery in problem-solving fosters curiosity and critical thinking, simulating the learning process for users.

Evaluating LLMs: The Technical Transparency Index

To facilitate systematic evaluation and comparison of model replication efforts, scholars have proposed the **Technical Transparency Index (TTI)**. This framework quantifies the transparency and reproducibility of implementations across key dimensions, offering a standardized approach to assessing foundational model performance and reliability.

Index 1: Data Transparency

This index evaluates the clarity and rigor with which datasets are documented and described, ensuring transparency in their origin, preparation, and application. Transparent data practices are crucial, particularly when datasets serve as seed data for generating synthetic long-thought datasets.

- **Data Source**: Examines whether the origins of the data are clearly specified, including detailed descriptions of datasets and their sources. It evaluates the explicit mention of dataset names, providers, or related publications.
- **Data Selection Process**: Focuses on the criteria and methodology used for filtering, cleaning, and preprocessing datasets before application in downstream tasks such as supervised fine-tuning (SFT), reinforcement learning (RL), or search algorithms.

Index 2: Methodology Transparency

Methodology transparency ensures that the techniques and processes employed are sufficiently detailed to enable independent reproduction and validation. This index evaluates various components of foundational models:

- **Foundation Model Details**: Assesses the depth of information provided about the base model, including architectural specifics (e.g., transformer layers, attention mechanisms) and parameter size.
- **Search Algorithm**: Evaluates the documentation of search algorithms used during inference, such as beam search or Monte Carlo Tree Search (MCTS).

Detailed descriptions of parameters, step-by-step processes, and any custom modifications are critical.

- **Reinforcement Learning (RL) Algorithm**: Reviews RL or preference learning methods, including details on reward functions, optimization goals, and training dynamics.
- **Long-Thought Synthetic Algorithm**: Examines processes for synthesizing long-thought datasets, including specific heuristics, rules, or algorithms used in data generation or selection.
- **Training Details**: Assesses the documentation of training procedures, covering key hyperparameters (e.g., learning rate, batch size, optimizer types) and the overall training configuration.
- **Effectiveness Validation**: Measures the rigor of validation processes for each method, ensuring empirical evidence supports claims about the importance of specific techniques.

Index 3: Evaluation Transparency

This index evaluates how clearly and comprehensively model performance is assessed:

- **Benchmark Usage**: Considers the appropriateness of benchmarks selected for evaluating task and domain-specific performance.
- **Evaluation Metrics**: Reviews the clarity and relevance of metrics used to quantify model performance, including any customizations introduced to address unique aspects of the evaluation.

Index 4: Open-Source Resources

Open-source contributions are integral to fostering reproducibility and enabling the research community to build upon existing work. This index evaluates the accessibility of critical resources:

- **Data**: Assesses whether post-training raw data and synthesized datasets (e.g., O1-like datasets) are publicly available, significantly enhancing reproducibility and enabling broader experimentation.
- **Model Weights**: Evaluates the availability of trained model weights, which facilitate replication and further optimization efforts.
- **Code**: Considers the comprehensiveness of released codebases, particularly whether they include scripts for both training and evaluation, along with thorough documentation.
- **Documentation**: Reviews supplemental materials, such as research papers, technical reports, or blog posts, to determine whether they clearly explain methodologies, results, and underlying ideas, providing actionable insights for researchers and practitioners.

Lessons Learnt

The trajectory of artificial intelligence (AI) research and technology faces a series of profound challenges that pose significant risks to the field’s long-term progress. These challenges span technical limitations, cultural shifts within the research community, and the erosion of foundational educational principles.

Surface Appeal

AI models can achieve rapid performance gains by leveraging sophisticated reasoning patterns with relatively straightforward implementations. This accessibility has driven widespread adoption, particularly among organizations eager to showcase capabilities quickly. However, the convenience of these methods often obscures their hidden costs, which could hinder the field’s growth and long-term sustainability.

Performance Ceiling

One of the most pressing technical concerns lies in the inherent limitations of distillation-based training approaches. Models trained through distillation are fundamentally constrained by the capabilities of their teacher models, creating a "ceiling effect."

Regardless of the sophistication of the distillation process, these models are unable to surpass the original teacher’s performance. This limitation is especially problematic when extending model capabilities to new domains or addressing previously unseen challenges, thereby restricting their potential for innovation.

Missed Opportunities for Innovation

The popularity of widely adopted training methods has diverted attention away from critical areas of technical innovation. True breakthroughs in AI are likely to emerge not only from solving complex problems but also from advancing capabilities in inference-time scaling and search optimization. However, reliance on distillation often bypasses these foundational challenges. This trend risks creating a widening technological gap between organizations that develop core technologies and those that depend primarily on distillation. As this gap grows, bridging it may become increasingly difficult, further polarizing the AI research landscape.

Shift in Research Culture

The ease of achieving “quick wins” has begun to reshape the culture of AI research. Efforts to tackle fundamental challenges are increasingly deprioritized in favor of faster, more accessible solutions. This shift has led to reduced investment in advanced computational infrastructure and diminished emphasis on developing sophisticated search and reasoning algorithms.

A self-reinforcing cycle emerges, where limited infrastructure restricts research possibilities, further encouraging reliance on simpler methods, and ultimately stifling innovation.

Erosion of Foundational Skills

The most alarming consequence of these trends lies in their impact on the education and training of future AI researchers. Widespread adoption of shortcut solutions risks eroding the ability of students and early-career researchers to engage with complex technical challenges from first principles. Fundamental scientific skills, such as problem-solving and systematic algorithm design, are at risk of being overshadowed by a focus on optimization and prompt engineering.

This transition from "understanding how it works" to "knowing what works" represents a significant shift in research mentality, with far-reaching implications for the field's innovative capacity.

Decay of First Principles Thinking

The decay of first-principles thinking undermines the very foundation of scientific innovation. The process of designing search algorithms, optimizing inference, and building reasoning mechanisms from scratch forces researchers to deeply engage with model behavior, limitations, and algorithmic intricacies. These experiences develop critical intuition and systematic problem-solving skills. Without such challenges, future researchers may become adept at applying existing solutions but lack the ability to develop new innovations from fundamental principles.

Impact on Academic Research

The implications of these trends extend to the broader academic environment. Universities, traditionally at the forefront of fundamental innovation, face mounting pressure to prioritize quick results over deeper technical investigations. Students may be discouraged from pursuing challenging research directions, and the focus on performance metrics risks producing a generation of researchers skilled in optimization but deficient in creative and exploratory capacity.

A Growing Divide

Over time, these factors are likely to exacerbate disparities within the AI research ecosystem. Organizations with the resources to develop foundational technologies, such as advanced search and inference mechanisms, will gain an increasingly significant advantage. Meanwhile, those reliant on distillation may remain confined to incremental improvements. This growing divide threatens to concentrate genuine breakthroughs within a small number of well-resourced entities, leaving the broader research community struggling to keep pace.

RECOMMENDATIONS FOR FUTURE RESEARCH

To confront the profound challenges facing the long-term development of AI technology and its research community, several key recommendations are proposed:

Organizations should adopt a balanced approach to their research portfolios, integrating both widely-used training methods and foundational research into search algorithms and inference optimization. While investment in advanced computing infrastructure remains critical, equal emphasis must be placed on developing core competencies in search and inference techniques to drive sustainable innovation. A comprehensive strategy that couples immediate performance gains with the pursuit of groundbreaking advancements will position organizations for long-term success.

In academia, the training of future AI researchers must be reimagined to cultivate both practical expertise and a deep understanding of fundamental principles. This requires:

1. **Balanced Curricula**: Educational programs should equally emphasize practical applications and theoretical foundations, ensuring that students are equipped with both technical skills and a strong grasp of first principles.
2. **Structured Research Projects**: Research initiatives should be designed to encourage students to engage deeply with the underlying mechanisms of AI systems, fostering a culture of exploration and innovation.
3. **Cultural Shift**: Academic environments must prioritize long-term innovation over short-term achievements, valuing the process of inquiry and discovery as much as tangible outcomes.

As the field progresses, maintaining a balance between immediate performance improvements and long-term development will be vital. By investing in both, we can ensure the continued advancement of AI capabilities while cultivating a new generation of researchers who prioritize innovation and critical thinking.

CONCLUSION

This paper examined the significant challenges these methods present, including performance ceiling effect, missed opportunities for fundamental innovation, changes in research culture and a decline in first-principles thinking among future researchers

It addressed these challenges by proposing:

1. Adoption of open science practices
2. Redefinition of scientific communication
3. Establishment of foundational principles for AI-driven scientific discovery
4. Promotion of responsible development in AI

While this approach may yield significant performance improvements, its widespread yet opaque application raises serious concerns about the future of AI research.

The ultimate mission lies in nurturing human minds capable of first-principles thinking. These individuals are the true architects of AI's future, and their ability to innovate will determine the trajectory of the field.

REFERENCES

1. Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. 2020. "Language Models Are Few-Shot Learners."
2. Campbell, Murray, A. Joseph Hoane, and Feng-hsiung Hsu. 2002. "Deep Blue." Artificial Intelligence 134, no. 1: 57–83.

3. Chen, Guoxin, Minpeng Liao, Chengxi Li, and Kai Fan. 2024. "Alphamath Almost Zero: Process Supervision Without Process." arXiv preprint abs/2405.03553.
4. Chern, Ethan, Haoyang Zou, Xuefeng Li, Jiewen Hu, Kehua Feng, Junlong Li, and Pengfei Liu. 2023. "Generative AI for Math: Abel." https://github.com/GAIR-NLP/abel.
5. Chowdhery, Aakanksha, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, et al. 2022. "PaLM: Scaling Language Modeling with Pathways."
6. Cobbe, Karl, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, et al. 2021. "Training Verifiers to Solve Math Word Problems." arXiv abs/2110.14168.
7. Deng, Yuntian, Kiran Prasad, Roland Fernandez, Paul Smolensky, Vishrav Chaudhary, and Stuart Shieber. 2023. "Implicit Chain of Thought Reasoning via Knowledge Distillation." arXiv preprint arXiv:2311.01460.
8. Dohmatob, Elvis, Yunzhen Feng, Pu Yang, Francois Charton, and Julia Kempe. 2024. "A Tale of Tails: Model Collapse as a Change of Scaling Laws."
9. Feng, Guhao, Bohang Zhang, Yuntian Gu, Haotian Ye, Di He, and Liwei Wang. 2024. "Towards Revealing the Mystery Behind Chain of Thought: A Theoretical Perspective." Advances in Neural Information Processing Systems 36.
10. Gerstgrasser, Matthias, Rylan Schaeffer, Apratim Dey, Rafael Rafailov, Henry Sleight, John Hughes, Tomasz Korbak, et al. 2024. "Is Model Collapse Inevitable? Breaking the Curse of Recursion by Accumulating Real and Synthetic Data."
11. Hao, Shibo, Yi Gu, Haotian Luo, Tianyang Liu, Xiyan Shao, Xinyuan Wang, Shuhua Xie, et al. 2024. "LLM Reasoners: New Evaluation, Library, and Analysis of Step-by-Step Reasoning with Large Language Models." arXiv preprint abs/2404.05221.
12. Huang, Jiaxin, Shixiang Shane Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. 2022. "Large Language Models Can Self-Improve." arXiv abs/2210.11610.
13. Huang, Zhen, Zengzhi Wang, Shijie Xia, Xuefeng Li, Haoyang Zou, Ruijie Xu, Run-Ze Fan, et al. 2024. "Olympicarena: Benchmarking Multi-Discipline Cognitive Reasoning for Superintelligent AI."
14. Kaplan, Jared, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. "Scaling Laws for Neural Language Models."
15. Li, Xian, Ping Yu, Chunting Zhou, Timo Schick, Omer Levy, Luke Zettlemoyer, Jason Weston, and Mike Lewis. 2024a. "Self-Alignment with Instruction Backtranslation." In The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7–11, 2024. OpenReview.net.
16. Li, Zhiyuan, Hong Liu, Denny Zhou, and Tengyu Ma. 2024b. "Chain of Thought Empowers Transformers to Solve Inherently Serial Problems." arXiv preprint abs/2402.12875.
17. Lightman, Hunter, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, et al. 2024. "Let's Verify Step by Step." In The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7–11, 2024. OpenReview.net.
18. Luo, Liangchen, Yinxiao Liu, Rosanne Liu, Samrat Phatale, Harsh Lara, Yunxuan Li, Lei Shu, Yun Zhu, et al. 2024. "Improve Mathematical Reasoning in Language Models by Automated Process Supervision." arXiv preprint abs/2406.06592

.