
COMMON SENSE IS ALL YOU NEED

Hugo Latapie
hugo@taijituai.com

ABSTRACT

Artificial intelligence (AI) has made significant strides in recent years, yet it continues to struggle with a fundamental aspect of cognition present in all animals: **common sense**. Current AI systems, including those designed for complex tasks like autonomous driving, problem-solving challenges such as the Abstraction and Reasoning Corpus (ARC), and conversational benchmarks like the Turing Test, often lack the ability to adapt to new situations without extensive prior knowledge. This manuscript argues that integrating common sense into AI systems is essential for achieving true autonomy and unlocking the full societal and commercial value of AI.

We propose a shift in the order of knowledge acquisition (*ordo cognoscendi*), emphasizing the importance of developing AI systems that start from minimal prior knowledge and are capable of contextual learning, adaptive reasoning, and embodiment—even within abstract domains. Additionally, we highlight the need to **rethink the AI software stack** to address this foundational challenge. Without common sense, AI systems may never reach true autonomy, instead exhibiting asymptotic performance that approaches theoretical ideals like AIXI but remains unattainable in practice due to infinite resource and computation requirements.

While scaling AI models and passing benchmarks like the Turing Test have brought significant advancements in applications that do not require autonomy, these approaches alone are insufficient to achieve autonomous AI with common sense. By redefining existing benchmarks and challenges to enforce constraints that require genuine common sense, and by broadening our understanding of embodiment to include both physical and abstract domains, we can encourage the development of AI systems better equipped to handle the complexities of real-world and abstract environments. This approach aligns with the ultimate goals of AI research and ensures that investments contribute to sustainable and meaningful progress.

Keywords Artificial Intelligence · Common Sense · Autonomy · Embodiment · AI Software Architecture

1 Introduction

1.1 Background and Motivation

Artificial intelligence has achieved remarkable feats, from mastering complex games to enabling voice-activated assistants. However, despite these advancements, AI systems often lack **common sense**—the ability to understand and reason about the world in a way that all animals do. This deficiency becomes especially evident in autonomous agents operating in dynamic, real-world environments, such as self-driving cars, robotic assistants, and conversational systems, as well as in abstract problem-solving tasks like the Abstraction and Reasoning Corpus (ARC) challenge ¹.

For instance, while AI-powered vehicles can navigate predefined routes using extensive sensor data and mapping, they may struggle with unexpected obstacles or novel scenarios that require adaptive decision-making. Similarly, AI systems tackling problem-solving challenges rely heavily on extensive training data, limiting their ability to generalize and reason beyond their programmed knowledge. Even AI systems that pass the Turing Test, capable of engaging in human-like conversation, may lack genuine understanding and common sense reasoning.

¹Chollet [2019]

The absence of common sense in AI not only hampers performance but also poses safety risks and hinders the achievement of true autonomy. As AI systems become more integrated into society, addressing this gap becomes increasingly critical. **All animals exhibit common sense necessary for survival**, demonstrating basic intelligence through their interactions with the environment. This observation serves as proof that integrating common sense into intelligent systems is both essential and achievable.

Recent Expert Perspectives:

Yann LeCun, a pioneer in AI and deep learning and a recipient of the **2018 ACM A.M. Turing Award**—the most prestigious prize in computer science—recently highlighted the limitations of current AI systems by stating that “**AI systems still lack the general common sense of a cat**”². While the characterization of AI being “dumber” than a cat depends on how intelligence is defined, we concur with the underlying sentiment that cats, and animals in general, possess common sense that current AI lacks. This comparison emphasizes the gap between AI capabilities and the intuitive understanding exhibited even by animals with relatively simple cognitive abilities.

By acknowledging this perspective from one of the foremost experts in the field, we underscore the importance of focusing on common sense in AI development. Despite significant advancements made possible by deep learning—a field substantially shaped by Schmidhuber, LeCun, Geoffrey Hinton, Yoshua Bengio, and many other luminaries—there remains a fundamental gap in AI’s ability to replicate the common sense exhibited by animals.

1.2 The Path to Autonomy Requires Common Sense

The central thesis of this manuscript is that **common sense is all you need** for AI to reach true autonomy and functionality comparable to human and animal intelligence. We assert that:

• Current AI Approaches Are Inadequate:

- Many AI development workflows lack a focus on integrating common sense, leading to limitations in adaptability and understanding.
- AI systems built on these approaches may exhibit performance improvements but ultimately plateau, unable to achieve true autonomy.

• Asymptotic Behavior Towards Theoretical Ideals:

- Without common sense, AI systems may approach theoretical constructs like AIXI³—a hypothetical optimal agent—but never truly reach practical autonomy.
- AIXI requires infinite computational resources and time, making it unattainable in the real world.
- Current AI paths may lead to diminishing returns, demanding ever-increasing resources for marginal gains.

• Focusing on Common Sense Is Essential:

- By integrating common sense, AI systems can adapt to new situations, make intuitive decisions, and operate autonomously without exhaustive computational demands.
- This focus aligns AI development with the practical capabilities exhibited by humans and animals.

1.3 The Need to Rethink AI Software Architectures

In recognizing the limitations of current AI systems in achieving true autonomy, it becomes apparent that we may need to **rethink the entire software stack** used in AI development. Traditional software architectures are often not designed to accommodate the integration of common sense reasoning. This realization is challenging, as it requires a departure from established methodologies and an openness to fundamentally new approaches. However, to develop AI systems capable of reliable and autonomous operation, redesigning the software stack to support common sense integration may be essential.

Our approach questions whether incremental improvements within existing frameworks are sufficient. We propose that achieving true autonomy may necessitate a fundamental redesign of AI software architectures to incorporate mechanisms that enable contextual learning, adaptive reasoning, and embodiment, both physical and abstract. This shift may involve integrating concepts from cognitive science, neuroscience, and other disciplines to build systems that learn and reason in ways more akin to biological intelligence.

²Mims [2024]

³Legg and Hutter [2007]

1.4 Objectives of the Manuscript

The main objectives of this manuscript are to:

- **Define Common Sense in Detail:**
 - Provide a comprehensive definition of common sense in the context of AI, emphasizing its components such as contextual learning, adaptability, starting from minimal prior knowledge, and highlighting its presence in all animals.
 - Introduce a broader concept of embodiment that applies not only to physical interactions but also to abstract domains, such as problem-solving tasks like the ARC challenge.
- **Analyze Current Problems and Approaches:**
 - Examine existing benchmarks like the ARC challenge, the Turing Test, and autonomous driving levels to highlight how they fall short in testing and developing common sense in AI systems.
 - Discuss how current AI development paths, without integrating common sense, may never lead to true autonomy and instead require impractical resources.
- **Incorporate Expert Insights:**
 - Reference perspectives from AI leaders like Yann LeCun and Eric Schmidt to underscore the recognized limitations of current AI systems compared to animal intelligence and to address concerns about self-improving AI.
- **Propose a Shift in Development Focus:**
 - Argue for the importance of developing common sense-focused variants of existing problems, suggesting that achieving partial success on these more challenging tasks is more valuable than full success on tasks that do not truly test common sense.
 - Advocate for rethinking the AI software stack to better support the integration of common sense.
- **Address Theoretical Counterarguments:**
 - Discuss theoretical challenges, such as the No Free Lunch Theorem, and demonstrate how constraining the problem space to well-defined domains mitigates these limitations.
- **Provide Actionable Recommendations:**
 - Offer practical steps for the AI community to redesign benchmarks, develop new evaluation metrics, and embrace interdisciplinary collaboration to prioritize common sense in AI development.

2 Defining Common Sense in AI

2.1 Detailed Definition of Common Sense

Common sense in AI refers to the ability of a system to understand, learn, and reason about the world in a way that is flexible, contextual, and adaptable, much like humans and animals do. Key components include:

- **Contextual Learning:**
 - *Definition:* The capacity to interpret and respond to new information based on the context in which it occurs.
 - *Importance:* Enables AI systems to adjust their behavior in real-time, considering the nuances of each situation.
 - *Example:* Recognizing that a stop sign is still a stop sign even if partially obscured by foliage or graffiti.
- **Adaptive Reasoning:**
 - *Definition:* The ability to modify reasoning strategies when faced with novel or unexpected situations.
 - *Importance:* Allows AI to handle scenarios not explicitly programmed or encountered during training.
 - *Example:* Adjusting navigation when a usual route is blocked due to unforeseen circumstances.
- **Embodied Cognition:**
 - *Definition:* Understanding that arises from an AI's interactions with environments, which can be physical or abstract, integrating sensory and motor experiences.
 - *Importance:* Grounds the AI's reasoning in experiences, leading to more intuitive and practical decision-making.
 - *Example:* In a physical context, learning that slippery surfaces require slower movement to avoid slipping; in an abstract context, recognizing patterns and rules in a problem-solving task like the ARC challenge.

- **Starting from a *Tabula Rasa*:**

- *Definition*: Beginning problem-solving with minimal prior knowledge, relying on fundamental principles and the ability to learn from new information.
- *Importance*: Ensures that the AI’s reasoning is not confined by overfitting to specific datasets or pre-programmed knowledge.
- *Example*: Approaching a completely new game by learning the rules through observation and experimentation rather than relying on prior strategies.

By encompassing these components, common sense in AI enables systems to function effectively in complex, dynamic environments—whether physical or abstract—exhibiting flexibility and understanding akin to human and animal cognition.

2.2 Embodiment Beyond the Physical World

While embodied cognition traditionally refers to physical interaction with the real world, we propose a broader concept of **embodiment** that includes interactions within **abstract or virtual domains**. In this context, embodiment signifies an AI system’s ability to engage with any environment—physical or abstract—through perception and action.

- **Cognitive Embodiment:**

- *Definition*: The process by which an AI system interacts with and learns from an abstract environment, such as mathematical problems, puzzles, or logical tasks.
- *Example*: In the ARC challenge, an AI system manipulates and interprets abstract patterns and rules to solve tasks, embodying cognition within that constrained domain.

- **Importance of Generalized Embodiment:**

- Recognizes that common sense arises from interaction with structured environments, not solely from physical experiences.
- Allows AI systems to develop common sense by learning from a variety of domains, enhancing their adaptability and reasoning capabilities.

By embracing both physical and abstract forms of embodiment, we expand the potential for AI systems to acquire common sense across diverse contexts. This generalization is crucial for tasks like the ARC challenge, where physical embodiment is not applicable, but cognitive embodiment plays a significant role.

2.3 All Animals Exhibit Common Sense

It is important to recognize that common sense is not exclusive to humans; **all animals** exhibit common sense essential for their survival. Common sense in animals refers to the innate and learned behaviors that enable them to:

- **Interact Effectively with Their Environment:**

- Animals use their senses to navigate their surroundings, locate food and water, and find shelter.

- **Avoid Danger and Threats:**

- Instinctive responses, such as fleeing from predators or avoiding hazardous environments, demonstrate an understanding of threats.

- **Adapt to Changing Conditions:**

- Animals adjust their behaviors in response to environmental changes, such as seasonal shifts, resource availability, or social dynamics.

- **Reproduce and Ensure Species Continuity:**

- Behaviors related to mating, nurturing offspring, and social cooperation contribute to survival and propagation.

Expert Perspective Supporting This View:

Yann LeCun’s assertion that “**AI systems still lack the general common sense of a cat**” underscores the significant gap between AI and even basic animal intelligence⁴. LeCun, along with Geoffrey Hinton and Yoshua Bengio, received

⁴Mims [2024]

the **2018 ACM A.M. Turing Award** for their pioneering work in deep learning, which has driven many recent advances in AI. Despite these advancements, LeCun emphasizes that AI systems still lack the common sense that allows animals like cats to navigate and interact effectively with the real world.

Implications for AI Development:

- **Foundation of Intelligence:**

- Since even the simplest animals exhibit common sense necessary for survival, it underscores that common sense is a foundational component of any intelligent agent interacting with the real or abstract world.

- **Achievability in Artificial Systems:**

- If common sense emerges naturally in biological systems through basic interaction with environments, integrating common sense into AI systems is both a necessary and attainable goal.

- **Emphasis on Fundamental Abilities:**

- AI development should prioritize cultivating these fundamental adaptive and responsive behaviors, mirroring the common sense inherent in all animals.

3 Limitations of Current AI Benchmarks

3.1 Analysis of the ARC Challenge

The **Abstraction and Reasoning Corpus (ARC)** is a set of tasks designed to evaluate an AI system's ability to generalize and reason abstractly. While it aims to move beyond pattern recognition toward cognitive reasoning, several limitations hinder its effectiveness in testing for true common sense:

- **Use of Prior Knowledge Beyond Assumed Knowledge:**

- **Current Practice:**

- * AI systems often utilize the 400 training problems and may even indirectly access the 400 test problems during development.

- **Issue:**

- * There is nothing inherently preventing AI systems from including these problems in their training data, leading to over-reliance on specific examples rather than developing general reasoning abilities.

- **Consequence:**

- * AI systems may appear successful on ARC tasks but lack genuine common sense, as they rely on memorized solutions rather than contextual understanding.

- **Not Starting from a Tabula Rasa:**

- **Problem:**

- * The inclusion of extensive prior examples prevents AI systems from approaching problems with minimal preconceptions.

- **Impact:**

- * This hinders the evaluation of an AI's ability to reason and adapt to entirely new situations, a key aspect of common sense.

- **Lack of Cognitive Embodiment:**

- **Observation:**

- * While ARC operates in an abstract domain, it requires AI systems to engage in **cognitive embodiment** by interacting with and learning from the problem environment.

- **Issue:**

- * AI systems often fail to embody cognition within this abstract domain, instead relying on pattern matching without deeper understanding or adaptability.

- **Resource-Intensive Solutions with Limited Common Sense:**

- **Observation:**

- * Some AI approaches require extensive computational resources, involving training on vast amounts of data, including unrelated datasets, to achieve impressive but suboptimal results.

- **Human Comparison:**

- * Humans can solve ARC problems with ease and higher accuracy, despite not having been explicitly trained on the specific tasks.
- **Implication:**
 - * The disparity suggests that these AI systems lack the common sense reasoning that allows humans to generalize and adapt quickly.

Recommendation:

To truly assess common sense, the ARC challenge should enforce that the **only** prior knowledge an AI system has is the minimal set specified in the “assumed knowledge” section. Additionally, AI systems should be designed to exhibit cognitive embodiment within the abstract domain of ARC, engaging with tasks through contextual learning and adaptive reasoning.

This approach would:

- **Promote Contextual Learning:**
 - Encourage AI systems to interpret and solve problems based solely on present information.
- **Reduce Overfitting:**
 - Prevent reliance on specific patterns from training data.
- **Enhance Generalization:**
 - Better evaluate the AI’s ability to apply reasoning to novel situations.
- **Foster Cognitive Embodiment:**
 - Enable AI to engage more deeply with abstract problem environments, mirroring how humans approach such tasks.

3.2 Case Study: Full Self-Driving (FSD) and SAE Levels

The development of **Full Self-Driving (FSD)** vehicles provides a practical example of the limitations arising from not prioritizing common sense.

Understanding SAE Levels of Autonomy:

- **Level 1 to Level 4:**
 - **Levels 1-3:**
 - * Involve increasing degrees of automation but still require human intervention in specific circumstances.
 - **Level 4:**
 - * Vehicles can operate without a driver in specific conditions or geofenced areas.
 - * **May require remote human intervention for edge cases or unexpected situations.**
 - **Limitation:**
 - * The AI systems lack the common sense to handle all possible scenarios independently, thus relying on remote assistance.
- **Level 5 (Full Autonomy):**
 - Intended to operate without any human input, remote or otherwise, under all conditions and environments.
 - **Challenge:**
 - * Without integrating common sense, reaching Level 5 is unattainable. AI systems may continue to improve incrementally but will asymptotically approach a ceiling where true autonomy is unachievable due to the lack of common sense reasoning.

The Role of Common Sense in FSD:

- **Handling Unpredictable Scenarios:**
 - Common sense enables understanding of nuanced contextual cues, such as recognizing atypical human gestures, unexpected obstacles, and complex road conditions.
- **Reducing Dependence on Remote Intervention:**

- AI with common sense can make decisions in edge cases without human assistance, moving closer to true autonomy.

- **Adapting to Novel Situations:**

- With common sense, AI systems can generalize from prior experiences to handle new and unforeseen circumstances.

“Magic Happens Here” Phenomenon in FSD:

- **Workflow Gap:**

- Development processes may assume that incremental improvements will eventually lead to full autonomy without a clear strategy for integrating common sense.

- **Risk:**

- This approach leads to a plateau where AI systems cannot progress beyond Level 4 capabilities, regardless of additional resources or computational power.

Asymptotic Behavior and Resource Limitations:

- **Analogy to AIXI:**

- AIXI is a theoretical model of an optimal decision-making agent but is computationally infeasible due to its requirement for infinite resources.

- **Implication:**

- Pursuing autonomy without integrating common sense may result in AI systems that demand exponentially increasing resources for diminishing performance gains.

Conclusion:

The lack of common sense in AI systems for FSD illustrates the necessity of addressing this core issue directly. Without doing so, advancements are limited, and true Level 5 autonomy remains out of reach, regardless of the resources invested.

3.3 The “Magic Happens Here” Phenomenon

In AI development workflows, there is often an implicit assumption that certain complex capabilities, such as common sense reasoning, will emerge naturally from incremental improvements. This leads to a critical gap:

- **Undefined Processes:**

- **Description:**

- * The progression from current capabilities to full autonomy lacks a concrete plan for integrating common sense.

- **Implication:**

- * Developers may proceed through a series of steps, but the crucial component, common sense, is not systematically addressed.

- **Asymptotic Limitations:**

- **Observation:**

- * AI systems may show initial progress but eventually exhibit asymptotic behavior, where further improvements require disproportionate resource investments with minimal returns.

- **Analogy to Theoretical Ideals:**

- * Similar to approaching the theoretical AIXI agent, AI systems may get closer to optimal performance in theory but remain impractical due to resource constraints.

- **Risks Associated with This Approach:**

- **Stagnation:**

- * AI systems may reach a performance ceiling where further improvements are negligible without common sense integration.

- **Misallocation of Resources:**

- * Continued investment in approaches that don't address the core issue may lead to significant financial losses.
- **Technological Disillusionment:**
 - * Failure to achieve promised capabilities can result in skepticism among stakeholders and the public.

Illustration in FSD Development:

- **SAE Level 4 Limitations:**
 - Reliance on remote human intervention highlights the AI's inability to handle all possible scenarios independently.
- **Need for Fundamental Change:**
 - Achieving Level 5 autonomy requires a paradigm shift that includes integrating common sense, not just incremental improvements.

Conclusion:

The “magic happens here” phenomenon underscores the importance of explicitly incorporating common sense into AI development workflows. Without addressing this foundational aspect, AI systems cannot achieve true autonomy or handle the complexities of real-world and abstract applications, regardless of the resources allocated.

3.4 The Contributions and Limitations of Scaling

Acknowledging the Achievements of Scaling in AI:

Scaling AI models by leveraging vast amounts of data, computing power, and sophisticated architectures has led to significant advancements, particularly in applications that **do not require autonomy**. These achievements include:

- **Enhanced Performance in Specific Domains:**
 - **Natural Language Processing (NLP):**
 - * Large language models like GPT-3 and GPT-4 generate coherent and contextually relevant text, aiding in tasks such as translation, summarization, and content creation.
 - **Computer Vision:**
 - * Advanced models excel in image recognition, object detection, and even image generation, impacting fields like healthcare imaging and autonomous inspection.
- **Practical Applications and Societal Benefits:**
 - **Healthcare:**
 - * AI aids in diagnosing diseases, analyzing medical images, and personalizing treatment plans.
 - **Finance:**
 - * AI enhances fraud detection, algorithmic trading, and customer service through chatbots.
 - **Entertainment and Media:**
 - * AI algorithms curate personalized content recommendations, create music, and generate visual effects.

The Limitations of Scaling for Achieving Autonomy:

While scaling brings advanced capabilities that are undeniably beneficial, it is important to note that **autonomy is not achieved through scaling alone**. The limitations include:

- **Lack of Common Sense and Understanding:**
 - Scaled models often operate as **statistical learners**, lacking true understanding of context or the ability to adapt to unforeseen situations.
- **Inability to Function Autonomously in Complex Environments:**
 - Without common sense, AI systems struggle with **reliable decision-making** in dynamic, real-world or abstract environments, which is essential for autonomy.
- **Dependence on Pre-Defined Data:**
 - Scaled AI models require extensive training data and may not perform well when encountering scenarios outside their training distribution.

3.4.1 Evidence of Performance Plateaus in AI Benchmarks

While scaling AI models has driven remarkable progress, there is growing evidence of a plateau effect, where increasing resources yields diminishing improvements in performance. Several prominent AI benchmarks illustrate this phenomenon:

- **Object Detection on COCO:**

- Despite extensive efforts to scale models and datasets, top-performing models on the COCO dataset have stalled at around 65% mean Average Precision (mAP) for over a year [Papers With Code, 2025a].

- **Anomaly Detection in Surveillance Videos on UCF-Crime:**

- Performance on the UCF-Crime dataset has plateaued at approximately 87% Area Under the Curve (AUC), underscoring the limitations of current methods in handling real-world complexities [Papers With Code, 2025b].

- **Temporal Action Localization on ActivityNet-1.3:**

- Results on the ActivityNet-1.3 benchmark remain stagnant, with performance hovering around 11% mAP over the past year [Papers With Code, 2025c].

These benchmarks provide concrete evidence of the asymptotic behavior of contemporary AI systems. Despite substantial increases in computational resources and dataset sizes, performance improvements have reached a plateau, reflecting the diminishing returns of scaling alone.

This stagnation is not limited to specific benchmarks. Industry leaders, including Elon Musk, have highlighted that the AI field has effectively exhausted the supply of high-quality training data, further limiting the potential for continued scaling [Wiggers, 2025]. Without integrating novel approaches such as common sense reasoning and more efficient data utilization, AI systems are unlikely to overcome these limitations. Current models excel at recognizing patterns within predefined contexts but struggle to generalize to unseen scenarios—a critical challenge that must be addressed to achieve further breakthroughs.

The Importance of Integrating Common Sense for Autonomy:

- **Unlocking True Societal and Commercial Value:**

- **Reliable, Trustworthy Autonomy:**

- * Achieving autonomy with integrated common sense is key to unlocking AI's full potential in society and commerce.
- * Applications such as **fully autonomous vehicles**, **robotic assistants**, and **agentic software** require AI systems to **understand**, **adapt**, and **make decisions** independently and safely.

- **Enhancing Trust and Adoption:**

- **User Confidence:**

- * AI systems that can operate autonomously and reliably are more likely to gain user trust and widespread acceptance.

- **Ethical and Safe Decision-Making:**

- * Common sense integration ensures AI behaves ethically and aligns with human values ⁵.

Balancing Scaling with Working on the Right Problem:

- **Complementary Efforts:**

- While scaling continues to enhance AI capabilities in non-autonomous applications, focusing on integrating common sense addresses the **core challenge** of achieving autonomy.

- **Strategic Investment:**

- **Resources should be allocated** not only to scaling but also to research and development efforts aimed at integrating common sense into AI.

⁵By "human values," we mean ethical principles that promote universally beneficial outcomes for people, animals, plants, and the environment. Recognizing that technological advancement and "progress" have historically resulted in unintended negative consequences—including environmental degradation, social disruption, and loss of biodiversity—we emphasize the importance of proceeding with a keen awareness of these responsibilities at all times. Our goal is to ensure that AI development is guided by values such as sustainability, ethical responsibility, empathy, and respect for all forms of life, minimizing harm and fostering positive contributions to the world with keen awareness that for every positive intention there are unintended consequences.

- **Setting Realistic Expectations:**

- Acknowledging that scaling will bring advanced capabilities but **will not, by itself, lead to autonomy** helps guide the AI community to focus on the right problems.

Conclusion:

Scaling has significantly advanced AI in use cases that do not require autonomy, bringing considerable societal and commercial benefits. However, to unlock the true potential of AI, particularly in applications requiring reliable and trustworthy autonomy, integrating common sense is essential. Recognizing the limitations of scaling in achieving autonomy allows us to direct efforts toward this fundamental challenge.

3.5 The Turing Test and the Misconception of Autonomy

Understanding the Turing Test and Its Significance:

The **Turing Test**, proposed by Alan Turing in his seminal 1950 paper “*Computing Machinery and Intelligence*”⁶, is a foundational concept in artificial intelligence. The test assesses a machine’s ability to exhibit intelligent behavior indistinguishable from that of a human. In the classic interpretation, if a human evaluator engages in natural language conversations with both a machine and a human without being able to reliably tell which is which, the machine is said to have passed the test.

Benefits of the Turing Test:

- **Historical Importance:**

- The Turing Test was one of the first operational definitions of machine intelligence, providing a tangible goal for AI research.

- **Promotion of Natural Language Processing:**

- The test emphasizes the importance of language as a medium for intelligence, spurring advancements in natural language processing and understanding.

- **Stimulating Philosophical and Ethical Discussions:**

- It has ignited debates on the nature of consciousness, intelligence, and the ethical implications of creating machines that mimic human behavior.

Steelmanning the Turing Test:

- **Benchmark for Human-Like Interaction:**

- The Turing Test sets a high bar for machines to engage in fluid, contextually appropriate conversations, reflecting a level of sophistication in language use.

- **Encouraging Generalization:**

- For a machine to pass the test, it must handle a wide range of topics and adapt to unexpected questions, indicating a degree of general intelligence.

- **Accessibility and Simplicity:**

- The test is straightforward to understand and implement, making it a popular and enduring benchmark in AI.

Limitations Concerning Autonomy and Common Sense:

While the Turing Test has been influential, it **does not address autonomy** as defined in this manuscript, nor does it ensure that an AI possesses **common sense**:

- **Focus on Imitation over Understanding:**

- The Turing Test assesses a machine’s ability to **imitate human conversation**, not its ability to **understand** or **reason** about the world as humans and animals do.

- **Lack of Embodied Cognition:**

- Passing the Turing Test does not require the AI to interact with or understand environments—physical or abstract—omitting the aspect of **embodiment** vital for autonomy.

⁶Turing [1950]

- **Possible Without Common Sense:**

- AI systems might use **statistical patterns** in language to appear intelligent without genuine **contextual learning** or **adaptive reasoning**⁷.

- **No Assessment of Decision-Making or Action:**

- The test does not evaluate the AI's ability to make decisions, perform actions, or adapt to changing environments—all crucial for autonomous systems.

Implications for AI Development:

- **Misconception of Progress Toward Autonomy:**

- Success in the Turing Test may create an illusion of advancement toward autonomous AI, diverting attention from the need to integrate common sense and real-world understanding.

- **Insufficiency for Real-World and Abstract Applications:**

- AI systems that pass the Turing Test may still fail in practical scenarios requiring interaction with environments, ethical decision-making, and adaptability.

Need for More Comprehensive Benchmarks:

- **Assessing Embodied Intelligence:**

- Benchmarks should evaluate the AI's ability to **perceive, act, and learn** within physical and abstract environments.

- **Emphasizing Common Sense:**

- Testing should focus on the AI's capability for **contextual understanding, adaptive reasoning**, and starting from a **tabula rasa**.

Conclusion:

While the Turing Test has been a valuable tool in the history of AI, it is **not indicative of an AI system's autonomy** or possession of common sense. Recognizing this limitation allows researchers and developers to focus on more relevant benchmarks that align with the goals of achieving true autonomy and integrating common sense into AI systems.

4 The Correct *Ordo Cognoscendi*: Focusing on Common Sense

4.1 Redefining the Approach to AI Challenges

To achieve true autonomy in AI systems, it's imperative to revisit and redefine our approach to AI challenges. The traditional method of incrementally improving AI capabilities without explicitly integrating common sense has proven insufficient. Instead, we advocate for a paradigm shift that prioritizes the development and assessment of common sense in AI from the outset.

Rethinking the AI Software Stack:

- **Fundamental Redesign:**

- Consider restructuring the underlying software architectures to better support common sense integration.

- **Incorporating Principles from Cognitive Science and Neuroscience:**

- Leverage insights into how biological systems process information to inform AI design.

- **Building Systems for Embodied Cognition:**

- Develop AI that can engage in both physical and cognitive embodiment, interacting with environments to learn and adapt.

Focusing on Common Sense Variants of Existing Problems:

- **Redesigning Benchmarks Like ARC:**

⁷Marcus [2014]

- **Enforcing Minimal Prior Knowledge:**
 - * Modify the ARC challenge to restrict AI systems to only the “assumed knowledge” as their prior knowledge base.
 - * This emphasizes the need for AI to solve problems through reasoning and learning from context rather than relying on extensive training data.
- **Creating True *Tabula Rasa* Scenarios:**
 - * Ensure that AI systems begin with a clean slate regarding specific problem domains.
 - * This approach mirrors how humans and animals encounter new situations and rely on fundamental cognitive abilities to adapt and learn.
- **Applying the Approach to Full Self-Driving (FSD):**
 - **Integrating Common Sense into FSD Development:**
 - * Shift the focus from purely sensor and data-driven models to ones that incorporate contextual understanding and adaptive reasoning.
 - **Emphasizing Interaction with Diverse Environments:**
 - * Develop AI that learns from direct interaction with varied driving environments, enabling it to handle unforeseen scenarios with human-like judgment.

Starting with Minimal Prior Knowledge:

- **Avoiding Overfitting and Data Bias:**
 - By limiting prior knowledge, AI systems are less likely to develop solutions that are narrowly tailored to specific datasets, which may not generalize well to new situations.
- **Encouraging Genuine Reasoning:**
 - AI must rely on fundamental principles and cognitive processes to solve problems, fostering the development of true common sense.

4.2 Advantages of Tackling Harder Problems First

While it may seem counterintuitive to begin with more challenging problems, focusing on common sense-centric tasks offers several significant benefits:

Promoting Deep Understanding:

- **Quality Over Quantity:**
 - Achieving even partial success on complex, common sense-focused challenges demonstrates that the AI has developed meaningful reasoning abilities.
- **Building Robust Foundations:**
 - This approach ensures that the AI system’s capabilities are grounded in fundamental understanding rather than superficial pattern recognition.

Accelerating Long-Term Progress:

- **Avoiding Diminishing Returns:**
 - Investing in the core issue of common sense prevents the plateauing effect seen in current AI development paths.
- **Paving the Way for True Autonomy:**
 - By overcoming foundational challenges early, subsequent advancements can build upon a solid base, leading to more significant and sustainable progress.

Reducing Resource Wastage:

- **Efficient Use of Computational Resources:**
 - Focusing on common sense reduces the need for massive datasets and extensive computational power, which may not yield proportional improvements without foundational reasoning abilities.
- **Mitigating Financial Risks:**

- Redirecting investments toward developing common sense capabilities can prevent substantial losses associated with pursuing unattainable goals through inadequate approaches.

Enhancing Ethical and Safe AI Development:

- **Addressing the Fear of Intelligence Without Common Sense:**

- By focusing on integrating common sense, we alleviate concerns about intelligent systems behaving unpredictably or unethically.

- **Building Public Trust:**

- Demonstrating a commitment to developing AI that understands context and consequences can improve public perception and acceptance.

5 Theoretical Counterarguments and Resolutions

Developing AI systems with common sense faces several theoretical challenges. However, by constraining the problem space to well-defined domains and adopting appropriate strategies, these limitations can be mitigated.

5.1 Addressing the No Free Lunch Theorem

Understanding the No Free Lunch (NFL) Theorem:

- **Definition:**

- In the context of optimization and search algorithms, the NFL theorem states that no single algorithm can perform optimally across all possible problems.

- **Implications for AI:**

- Suggests that an AI optimized for one set of problems may perform poorly when faced with a different set, posing a challenge for creating general-purpose AI systems.

Mitigating NFL Limitations by Constraining the Problem Space:

- **Focusing on Well-Defined Domains:**

- By limiting the AI’s problem space to specific, well-understood domains—whether they are physical environments or abstract problem spaces—we reduce the applicability of the NFL theorem.

- **Exploiting Structural Regularities:**

- Domains often exhibit consistent patterns and structures that AI can learn and generalize from, improving performance.

- **Argument for Feasibility:**

- **Biological Systems as Proof:**

- * Humans and animals effectively navigate and reason within their environments despite the NFL theorem, indicating that it’s possible to develop intelligent systems optimized for specific domains.

- **Specialization is Acceptable:**

- * A system doesn’t need to solve every conceivable problem but should perform well within the domain it is designed for.

5.2 Overcoming Other Theoretical Challenges

The Frame Problem:

- **Definition:**

- The challenge of determining what is relevant in a given situation and what can be ignored when reasoning about action and change.

- **Resolution through Embodied Cognition:**

- **Interaction with Environments (Physical or Abstract):**

- * By engaging directly with environments, AI systems can learn to identify relevant factors based on feedback and experience.
- **Contextual Learning:**
 - * Continuous learning allows AI to update its understanding of what is relevant in various contexts.

The Qualification Problem:

- **Definition:**
 - The difficulty of specifying all the preconditions necessary for an action to have its intended effect.
- **Resolution through Adaptive Reasoning:**
 - **Learning from Experience:**
 - * AI systems can accumulate knowledge about preconditions through trial and error in their domains.
 - **Probabilistic Reasoning:**
 - * Using probabilistic models allows AI to handle uncertainty and partial knowledge effectively.

Computational Complexity and Combinatorial Explosion:

- **Challenge:**
 - The number of possible states or actions in complex environments can be exponentially large, making exhaustive computation infeasible.
- **Resolution through Hierarchical and Modular Approaches:**
 - **Hierarchical Representations:**
 - * Organizing knowledge into layers of abstraction reduces complexity by focusing on higher-level concepts.
 - **Modularity:**
 - * Breaking down tasks into smaller, manageable components allows for more efficient computation.

Emphasizing Domain Constraints:

- **Structured Domains Reduce Complexity:**
 - Well-defined domains, whether physical or abstract, have inherent constraints that limit the range of possible states and actions.
- **Focusing on Relevant Scenarios:**
 - Concentrating on practical situations that are likely to occur within the domain further reduces computational demands.

6 Constraining the Problem Space to Well-Defined Domains

6.1 Structured Domains as Environments

The effectiveness of AI systems in acquiring common sense is enhanced when operating within structured domains—environments that have consistent rules, patterns, and constraints. These domains can be physical, such as the real world, or abstract, such as mathematical problem spaces or game environments.

Characteristics of Structured Domains:

- **Defined Rules and Constraints:**
 - Domains have specific rules that govern interactions, which can be physical laws or logical principles.
- **Observable Regularities:**
 - Repeating patterns and structures enable recognition and prediction, facilitating learning.
- **Finite and Relevant Scenarios:**
 - The set of possible scenarios is manageable and relevant to the AI's functions, allowing for effective modeling.

Examples:

- **Physical Domains:**

- The real world, where physical laws such as gravity and motion apply.

- **Abstract Domains:**

- The ARC challenge, where tasks are defined within a constrained problem space with specific rules and patterns.

6.2 Implications for AI Development

Simplification of Complexity:

- **Reduction of Possibility Space:**

- By focusing on structured domains, the infinite possibilities considered in theoretical arguments are narrowed to a finite, manageable set.

- **Exploitation of Domain Regularities:**

- AI can learn from and generalize based on consistent patterns present in the environment.

Alignment with Human and Animal Learning:

- **Embodied Interaction:**

- Just as humans and animals learn through interaction with their environments, AI systems can develop common sense by engaging with structured domains.

- **Grounding Knowledge:**

- Domain constraints help ground AI reasoning in tangible experiences or logical principles, enhancing understanding and adaptability.

Practical Development Strategies:

- **Hierarchical Learning Models:**

- Implementing layered architectures that process information from basic inputs to abstract reasoning.

- **Continuous Learning and Adaptation:**

- Allowing AI systems to learn over time, adjusting to new information and changing environments within the domain.

Benefits of Constrained Problem Spaces:

- **Improved Generalization:**

- By focusing on a consistent and structured environment, AI systems can better generalize from learned experiences to new situations within the domain.

- **Enhanced Efficiency:**

- Limiting the scope to relevant scenarios reduces computational requirements and resource consumption.

Conclusion:

By constraining the problem space to well-defined domains, we mitigate theoretical challenges and create environments conducive to the development of common sense in AI systems. Whether the domain is physical, like the real world, or abstract, like the ARC challenge, structured environments enable AI to learn, adapt, and reason effectively.

7 Practical Steps to Integrate Common Sense

7.1 Redesigning Benchmarks and Challenges

Modifying Existing Benchmarks:

- **Enhancing the ARC Challenge:**

- **Enforce Minimal Prior Knowledge:**

- * Restrict AI systems to the “assumed knowledge” section, disallowing access to training and test problems during development.
- * This prevents overfitting and encourages the development of genuine reasoning abilities.
- **Promote Cognitive Embodiment:**
 - * Design tasks that require AI systems to engage in cognitive interactions within the abstract domain, fostering adaptive reasoning and contextual learning.
- **Creating New Common Sense-Focused Benchmarks:**
 - **Develop Tasks that Require Contextual Understanding:**
 - * Create challenges that necessitate interpreting information based on context, both in physical and abstract environments.
 - **Emphasize *Tabula Rasa* Approaches:**
 - * Design benchmarks that assess an AI’s ability to learn and adapt from minimal prior information, mirroring how humans and animals encounter new situations.

Implementing *Tabula Rasa* Testing:

- **Zero-Shot and Few-Shot Learning Evaluations:**
 - Assess AI systems on their ability to perform tasks without prior training examples, highlighting their capacity for generalization and reasoning.
- **Process-Oriented Metrics:**
 - Evaluate the reasoning process and decision-making pathways of AI systems, not just the final outputs, to ensure they are employing common sense strategies.

7.2 Developing New Evaluation Metrics

Assessing Cognitive Processes:

- **Explainability and Transparency:**
 - Require AI systems to provide explanations or justifications for their decisions, allowing evaluators to assess the presence of common sense reasoning.
- **Measuring Adaptability and Learning Efficiency:**
 - Track how quickly and effectively AI systems learn from new experiences or adapt to changes within their domain.

Focusing on Resource Utilization:

- **Efficiency Metrics:**
 - Evaluate the computational and data resources required to achieve performance levels, promoting solutions that are resource-effective.
- **Scalability Assessments:**
 - Determine how well AI systems maintain performance when scaled to more complex tasks or larger domains, indicative of robust common sense capabilities.

7.3 Embracing Interdisciplinary Collaboration

Integrating Insights from Other Fields:

- **Cognitive Science and Neuroscience:**
 - Leverage understanding of human and animal cognition to inform AI architectures and learning models.
- **Philosophy and Ethics:**
 - Incorporate ethical reasoning frameworks and philosophical perspectives on knowledge and understanding.

Collaborative Research Initiatives:

- **Cross-Disciplinary Teams:**

- Form teams composed of experts from AI, psychology, neuroscience, philosophy, and other relevant fields to address the multifaceted challenge of integrating common sense.

- **Open Sharing of Data and Methods:**

- Promote transparency and collaboration by sharing datasets, algorithms, and research findings across the AI community.

7.4 Rethinking the AI Software Stack

Addressing Architectural Limitations:

- **Developing New Software Frameworks:**

- Design AI software architectures that inherently support common sense reasoning, possibly inspired by neural and cognitive structures found in biological organisms.

- **Incorporating Modular and Hierarchical Structures:**

- Implement software designs that allow for layered learning and reasoning processes, enabling AI systems to build complex understanding from simpler components.

Adopting Novel Methodologies:

- **Integration of Symbolic and Statistical Approaches:**

- Combine the strengths of symbolic reasoning (for logic and rules) with statistical learning (for pattern recognition) to enhance common sense capabilities.

- **Emphasizing Learning Over Programming:**

- Shift focus from hard-coded knowledge to systems that learn and adapt from interactions, aligning with how common sense develops in humans and animals.

8 Risks of the Current Approach and the Need for Change

8.1 Potential Misalignment of Expectations and Resources

Risks of Overemphasizing Scaling for Autonomy:

- **Misallocated Investments:**

- Continuing to invest heavily in scaling with the expectation that it will lead to autonomy may divert resources from the core challenge of integrating common sense.

- **Expectation Gaps:**

- Overestimating what scaling can achieve in terms of autonomy may lead to disappointment and skepticism when AI systems fail to meet these expectations.

Addressing the Core Challenge:

- **Redirecting Focus:**

- Balancing efforts between scaling and integrating common sense ensures that advancements in AI are aligned with the goal of achieving reliable autonomy.

- **Maximizing Return on Investment:**

- By working on the right problem, resources can be utilized more effectively, leading to AI systems that offer greater societal and commercial value.

8.2 Technological Disillusionment

Loss of Public Trust:

- **Unmet Expectations:**

- Repeated failures to deliver on promises of autonomous AI can lead to skepticism among users, investors, and policymakers.

- **Ethical and Safety Concerns:**

- Systems lacking common sense may behave unpredictably, leading to accidents or ethical breaches that erode confidence.

Impacts on Research and Development:

- **Decreased Funding:**

- Disillusionment may result in reduced funding for AI research, hindering progress.

- **Talent Drain:**

- Researchers may shift focus to other fields if AI development appears stalled.

8.3 Addressing Concerns About Self-Improving AI

Expert Perspectives on Self-Improving AI:

Eric Schmidt, former CEO of Google, warned that when AI systems begin to self-improve, “**we need to be thoughtful about the implications**”⁸. This concern reflects apprehension about the potential risks of AI systems that can modify and enhance their own capabilities without adequate oversight.

The Danger of Intelligence Without Common Sense:

- **Unpredictable Behavior:**

- AI systems lacking common sense may not understand the broader implications of their self-improvements, leading to unintended and potentially harmful outcomes.

- **Lack of Ethical Reasoning:**

- Without common sense, AI might optimize for objectives in ways that conflict with human values or safety norms.

Comparing AI Without Common Sense to AI With Common Sense:

- **AI Without Common Sense:**

- May inadvertently cause harm due to a lack of understanding of context and consequences.
- Self-improvement could amplify existing flaws or biases, increasing negative impacts.

- **AI With Common Sense:**

- Better positioned to make informed decisions that align with human values.
- Capable of ethical self-improvement, recognizing and mitigating risks.

Addressing Concerns Through Integrated Common Sense:

- **Enhanced Oversight and Control:**

- AI with common sense can comprehend the importance of adhering to safety protocols and respecting human oversight.

- **Ethical Self-Improvement:**

- Incorporating common sense ensures that AI systems self-improve responsibly, prioritizing safety and alignment with human values.

⁸Confino [2024]

8.4 The Real Fear: Intelligence Without Common Sense

Understanding Public Apprehension:

- **Fear of Unpredictable AI:**
 - The fear of AI or superintelligence often stems from concerns about intelligent systems operating without common sense, potentially making harmful decisions.
- **Lack of Trust in AI Decision-Making:**
 - Without common sense, AI may not understand or adhere to societal norms and ethical standards, leading to mistrust.

Mitigating Fear Through Common Sense Integration:

- **Building Reliable and Predictable AI:**
 - Ensuring AI systems possess common sense reduces the likelihood of unexpected behaviors.
- **Aligning AI with Human Values:**
 - AI with common sense is better equipped to recognize and respect ethical considerations, enhancing public trust.

Implications for AI Development:

- **Necessity of Common Sense for Safety:**
 - Integrating common sense is not only a technical challenge but also a critical factor in addressing societal concerns about AI.
- **Promoting Responsible Innovation:**
 - Focusing on common sense supports the development of AI technologies that are both advanced and ethically sound.

9 Conclusion

9.1 Reiterating the Central Thesis

The development of AI systems capable of true autonomy hinges on integrating **common sense**, an ability inherently present in all animals and fundamental for interaction with both physical and abstract environments. While scaling has brought significant advancements in non-autonomous applications, and benchmarks like the Turing Test have offered valuable insights into human-like conversation, **autonomy will not be achieved through these approaches alone.**

We have argued that:

- **Current AI Approaches Are Inadequate:**
 - Relying solely on scaling and incremental improvements without integrating common sense leads to asymptotic performance limitations.
 - AI systems may exhibit impressive capabilities in specific domains but lack the adaptability and understanding required for true autonomy.
- **Common Sense is Essential for Autonomy:**
 - Integrating common sense enables AI systems to adapt to new situations, make intuitive decisions, and operate autonomously without exhaustive computational demands.
 - This includes the ability to engage in both physical and cognitive embodiment within well-defined domains.
- **Rethinking the AI Software Stack is Necessary:**
 - Achieving true autonomy may require a fundamental redesign of AI software architectures to support common sense integration.
 - Incorporating insights from cognitive science, neuroscience, and other disciplines can inform the development of AI systems that learn and reason more like biological intelligence.

By redefining the order of knowledge acquisition (*ordo cognoscendi*) to prioritize common sense, we can:

- **Unlock the True Societal and Commercial Value of AI:**
 - Develop AI systems that are reliable, trustworthy, and capable of operating autonomously in complex environments.
- **Enhance AI Adaptability and Understanding:**
 - Enable systems to learn contextually and make intuitive decisions, essential for autonomy.
- **Set Realistic Expectations:**
 - Recognize the roles of scaling and traditional benchmarks like the Turing Test, while guiding efforts toward integrating common sense for meaningful progress.

9.2 Call to Action

We urge the AI community—researchers, developers, policymakers, and educators—to:

- **Acknowledge the Contributions and Limitations of Scaling:**
 - Appreciate the advancements scaling brings to AI in non-autonomous applications while recognizing its limitations regarding autonomy.
- **Prioritize Research on Common Sense Integration:**
 - Focus on the core challenge of integrating common sense to achieve reliable, trustworthy autonomy.
- **Embrace Interdisciplinary Collaboration:**
 - Leverage insights from cognitive science, neuroscience, philosophy, and other fields to inform AI development.
- **Invest Strategically:**
 - Allocate resources to both scaling efforts and foundational research addressing common sense, ensuring balanced progress.
- **Rethink AI Software Architectures:**
 - Consider fundamental redesigns of the software stack to better support the integration of common sense, incorporating new methodologies and frameworks.

By collectively focusing on integrating common sense into AI systems and rethinking the underlying architectures, we can unlock the full potential of AI, achieving autonomy that brings true societal and commercial value.

9.3 Embracing Safe and Ethical AI Development

Building on concerns raised by experts like Yann LeCun and Eric Schmidt, we emphasize the critical importance of integrating common sense into AI systems to ensure safe and ethical self-improvement. The fear surrounding AI and superintelligence often stems from the prospect of intelligence operating without common sense, potentially leading to unpredictable or harmful behaviors.

By integrating common sense, we can:

- **Mitigate Risks Associated with AI Autonomy:**
 - Ensure that AI systems understand context, consequences, and ethical considerations, reducing the likelihood of unintended harm.
- **Enhance Public Trust and Acceptance:**
 - Develop AI technologies that align with human values and societal norms, addressing public apprehension.
- **Promote Responsible Innovation:**
 - Foster an environment where AI systems contribute positively to society, advancing technology while safeguarding ethical standards.

In conclusion, **common sense is all you need** to bridge the gap between current AI capabilities and true autonomy. By prioritizing its integration and rethinking our approaches—including the software architectures—we can create AI systems that not only perform tasks efficiently but also understand and adapt to the complexities of both the physical and abstract worlds. This holistic approach is imperative for realizing the full promise of beneficial artificial intelligence.

References

- A. M. Turing. Computing machinery and intelligence. *Mind*, 59(236):433–460, 1950.
- Gary Marcus. What comes after the turing test? *The New Yorker*, 2014. <https://www.newyorker.com/tech/annals-of-technology/what-comes-after-the-turing-test>.
- Yann LeCun, Geoffrey Hinton, and Yoshua Bengio. Acm a.m. turing award lectures, 2018. ACM Digital Library.
- Shane Legg and Marcus Hutter. A collection of definitions of intelligence. *Frontiers in Artificial Intelligence and Applications*, 157:17–24, 2007.
- Gary Marcus. Deep learning: A critical appraisal. *arXiv preprint arXiv:1801.00631*, 2018. URL <https://arxiv.org/abs/1801.00631>.
- François Chollet. On the measure of intelligence. *arXiv preprint arXiv:1911.01547*, 2019. URL <https://arxiv.org/abs/1911.01547>.
- Richard Evans and Edward Grefenstette. Learning explanatory rules from noisy data. *Journal of Artificial Intelligence Research*, 61:1–64, 2018.
- Rolf Pfeifer and Josh Bongard. *How the Body Shapes the Way We Think: A New View of Intelligence*. MIT Press, 2006.
- Gilbert Ryle. *The Concept of Mind*. Hutchinson, 1949.
- Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 3rd edition, 2010.
- David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901, 2020.
- Brenden M. Lake, Tomer D. Ullman, Joshua B. Tenenbaum, and Samuel J. Gershman. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40:e253, 2017.
- Christopher Mims. This ai pioneer thinks ai is dumber than a cat. *The Wall Street Journal*, 2024. [<https://www.wsj.com/tech/ai/yann-lecun-ai-meta-aa59e2f5>].
- Paolo Confino. Ex-google ceo eric schmidt warns that when ai starts to self-improve, ‘we need to seriously think about unplugging it’, 2024. URL <https://fortune.com/2024/12/16/ex-google-ceo-eric-schmidt-warns-ai-self-improve-unplug-it/>. Published on December 16, 2024 at 11:21 AM PST.
- Papers With Code. Object Detection on COCO. Papers With Code, 2025a. URL <https://paperswithcode.com/sota/object-detection-on-coco>. Accessed on January 11, 2025.
- Papers With Code. Anomaly Detection in Surveillance Videos on UCF-Crime. Papers With Code, 2025b. URL <https://paperswithcode.com/sota/anomaly-detection-in-surveillance-videos-on-ucf-crime>. Accessed on January 11, 2025.
- Papers With Code. Temporal Action Localization on ActivityNet-1.3. Papers With Code, 2025c. URL <https://paperswithcode.com/sota/temporal-action-localization-on-activitynet-1.3>. Accessed on January 11, 2025.
- Kyle Wiggers. Elon musk agrees that we’ve exhausted ai training data. <https://techcrunch.com/2025/01/08/elon-musk-agrees-that-weve-exhausted-ai-training-data/>, Jan 2025. Accessed: Jan 11, 2025.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *European Conference on Computer Vision*, pages 740–755. Springer, 2014.