2025/11/18 17:04 1/10 Links für 2023 KW 34

Links für 2023 KW 34

Zusammengefasst von LlongOrca.

Paper Review: Llama 2: Open Foundation and Fine-Tuned Chat Models:

Llama 2 is an open-source large language model (LLM) developed by the authors to surpass the performance of existing models on various benchmarks. The model offers a range of sizes from 7 billion to 70 billion parameters and is designed for dialogue applications, such as chatbots. The Llama 2-Chat fine-tuned versions have demonstrated better safety and helpfulness than existing open-source chat models in human evaluations.

The Llama 2 model family was created from the pretraining methodology of Llama and features enhancements like enhanced data cleaning, updated data mixes, and doubling the context length for improved performance. The group-query attention (GQA) is incorporated to boost the inference scalability for larger models.

The authors used a mixture of publicly accessible data to train Llama 2 models, excluding any data from Meta's products or services and removing information from sites known for storing personal information about private individuals. They trained the model on 2 trillion tokens of data and balanced performance with cost by up-sampling factual sources.

In fine-tuning, Llama 2 models were evaluated in terms of helpfulness and safety through human preferences. The authors developed a reward model that could be used to optimize the Llama 2-Chat model for better alignment with user requirements. Using reinforcement learning from human feedback (RLHF), the models were trained based on human preference data, and the results showcased improvements in performance over time.

Safety is a crucial aspect of these models; therefore, the authors introduced safety fine-tuning methods, including supervised safety fine-tuning, safety RLHF (reinforcement learning from human feedback), and safety context distillation to improve model robustness. The Llama 2-Chat models were evaluated on safety benchmarks such as TruthfulQA for truthfulness, ToxiGen for toxicity, and BOLD for bias.

The researchers also employed a red teaming approach to identify potential risks in the model, which involved more than 350 professionals from diverse fields evaluating different risk categories. The findings helped improve model safety and robustness through iterative improvements.

In summary, Llama 2 is an advanced open-source large language model that offers improved performance, helpfulness, and safety compared to existing models. Its fine-tuning techniques ensure better alignment with user requirements while maintaining a high level of safety and robustness.

Lie still in bed:

The article discusses how the author struggled with maintaining a regular sleep schedule after university and found an effective solution - lying still in bed at the target sleep time every day. The author claims that most people don't have problems falling asleep, but instead spend too much time on their phones or watching TV before going to sleep. By using this simple advice and practicing consistency, the author eventually became a morning person and applied it to other areas of life. The

article highlights the importance of willpower for achieving progress, acknowledging that small steps lead to significant improvements over time.

Python Type Hints - *args and **kwargs - Adam Johnson:

The article explains how to use type hints with Python's variable argument operators, * and (often called *args and kwargs). It provides an example of a function that takes both positional arguments and keyword arguments and details the proper way to define its type hints. The author clarifies that specifying only the contained argument types is enough; the type checker automatically adds the necessary container types, like Tuple[_, ...] or Dict[str, _].

puzzle.telegnom.org:		
TODO		
blog.gopenai.com:		
TODO		

☐ Python Bindings for Ilama.cpp:

Llama-cpp-python is a Python library that provides bindings for @ggerganov's llama.cpp, allowing low-level access to C API via ctypes interface and high-level Python API for text completion. The library also offers compatibility with OpenAI-like API and LangChain. Installation from PyPI is possible, with additional information on installation available at https://llama-cpp-python.readthedocs.io/en/latest/.

The Secret Sauce behind 100K context window in LLMs: all tricks in one place:

In summary, techniques to speed up training and inference of LLMs include using large context windows (up to 100K input tokens) during training and inference. These techniques include ALiBi positional embedding, Sparse Attention, FlashAttention, Multi-Query attention, Conditional computation, and utilizing 80GB A100 GPUs. By optimizing these processes, LLM models can effectively process longer input sequences for improved performance and accuracy in natural language processing tasks.

Big content update 1.0.7 for Idle Tower Builder:

Idle Tower Builder v.1.0.7 has been released with new machines, balance adjustments, and other improvements. It is now available on the App Store, Google Play, and Itch. New features include a Fountain, Gunpowder Factory, Cannon, Bungee Jumper, Moon stairs, Moon Tower, Zeppelin upgrades,

2025/11/18 17:04 3/10 Links für 2023 KW 34

and more. Users can expect increased resource capacity and automated facility purchasing after building restarts with the new version.

Raketenflieger Timmi:

Raketenflieger Timmi ist eine deutsche Zeichentrickserie für Kinder, die im Jahr 2014 erstmals ausgestrahlt wurde. Die Serie wird von MotionWorks produziert und handelt von Astronaut Timmi, der mit seiner selbstgebauten Rakete zusammen mit seinem Freund Teddynaut Teddy auf verschiedenen Planeten durch das All fliegt. Jede Folge hat eine Länge von rund vier Minuten und zeigt den Charakter des jungen Helden und seine Abenteuer im All. Die Serie wurde 2016 mit dem Kinderfernsehpreis EMIL ausgezeichnet.

Flatpak Is Not the Future:

I appreciate your thorough analysis of Flatpak as well as other packaging solutions. You raise several valid concerns about these technologies, such as their size, incompatibilities with system libraries, and complexity. While there are undoubtedly some advantages to sandboxed app stores, the drawbacks you outline should be carefully considered before implementing these solutions on a large scale.

The desire for compatibility and stability across various Linux distributions is critical to attracting proprietary software developers to the platform. As such, efforts should focus on developing better tools and strategies that allow apps to be ported easily between different systems while maintaining backwards compatibility. By addressing these concerns, we can create a more inviting environment for both open source and proprietary software alike.

Giraffe - Long Context LLMs - The Abacus.Al Blog:

The article discusses the development of Giraffe, a new family of models that are finetuned from base LLaMA and LLaMA2 for expanding context lengths in Large Language Models (LLMs). The paper explores the concept of context length extrapolation, which allows trained models to be applied to longer contexts without further training. It highlights the challenges in evaluating LLM performance using existing metric methods and introduces new tasks such as LongChat-Lines, FreeFormQA, and AlteredQA to assess accuracy and recall in these models. The paper aims to continue researching context length extrapolation of LLMs to resolve existing issues in the field.

www.datapen.io:

Free Online Resources for Data People

Your attention didn't collapse. It was stolen:

Last update: 2024/08/26 08:35

This article discusses the impact of technology and our modern lifestyle on attention span, focus, and mental health. The author recounts his personal journey with distractions in daily life, leading him to spend time in a tech-free environment and learning from experts in the field. He suggests that individual abstinence is not the solution; instead, systemic changes are needed to address environmental factors that contribute to our attention problems.

Training immune cells to remove 'trash' helps resolve lung inflammation:

Inflammation is a natural part of the immune system but can become hyperactivated in lungs leading to fatal outcomes. Researchers at the University of Illinois Chicago have found that lung cells called alveolar macrophages counteract inflammation by clearing debris and releasing anti-inflammatory proteins. After initial exposure to a bacterial toxin, these trained cells help reduce severity of inflammation in subsequent exposures. The findings suggest that these trained cells could be used as part of a cell therapy to prevent excessive inflammation.

Airlines are cracking down on this travel hack and one website is in the crosshairs:

Skiplagging, or hidden city travel, is a sneaky travel hack that involves purchasing multi-stop fares with a layover in the desired destination instead of booking a direct flight. Customers disembark at their stopover location, leaving an empty seat on the final leg of their journey, and save money in the process. However, airlines are cracking down on this practice, claiming it violates their rules and costs them revenue. Some analysts argue that skiplagging is a form of fraud because customers are purchasing flights they don't intend to travel on. The debate over skiplagging continues, but it remains popular among those seeking cheap flight options.

My hobby became my job, 50% extra pay, just needed to let go of GPLv3 | authentik:

This article discusses the challenges faced by an individual who initially started working on a project called authentik with an MIT license. However, when Elastic called out AWS for trademark abuse, the author switched to GPLv3 to prevent their software from facing a similar issue. Two years later, they decided to switch back to MIT as it offered better legal protection and was more business-friendly. This change of licenses sparked reflections on what it means to build a company around an open source project. The author concluded by stating that being a public benefit company has helped strike a balance between commercial viability and upholding the values of open source development.

Big Ass Data Broker Opt-Out List:

This article provides a comprehensive list of data brokers and methods to opt-out or remove personal information from various websites. It also mentions special circumstances, such as being a victim of crime or stalking, that may allow for additional opt-outs. The author recommends exploring privacy options available through state motor vehicle departments and contacting phone companies to opt out of sales of numbers. Some paid services for removing personal information are also mentioned, although the effectiveness is not guaranteed.

2025/11/18 17:04 5/10 Links für 2023 KW 34

Astrum - YouTube:

The article discusses the use of artificial intelligence (AI) in business, focusing on how it can assist with decision-making processes and automation. It highlights examples such as self-driving cars and digital marketing to demonstrate how AI is revolutionizing various industries.

OpenTF created a fork of Terraform!:

In August 2023, HashiCorp announced a license change for their core products, including Terraform, to the Business Source License (BSL). To preserve Terraform's open-source nature, the OpenTF manifesto was released, resulting in over 100 companies and hundreds of individuals pledging resources to maintain Terraform's open-source status. Since no license reversal occurred, OpenTF, a fork of Terraform, has been created. The project focuses on remaining true to its open-source values, community-driven governance, and vendor neutrality. It is in the process of joining the Linux Foundation with intentions to be part of the Cloud Native Computing Foundation.

FDA approving drugs after fewer trials, providing less information to public, studies find:

The U.S. Food and Drug Administration (FDA) is approving more novel pharmaceutical drugs based on single clinical trials, with less public disclosure about those trials than before, according to studies from Oregon State University. Researchers suggest that there is a need for greater transparency around how drugs receive approval, including showing all results or trials completed. The FDA has relaxed some standards since the 21st Century Cures Act passed in 2016, allowing treatments for priority health conditions to be approved with fewer supporting studies and placing less emphasis on randomized clinical trials in certain cases.

☐ Improving sleep:

The article discusses the author's experience with sleeping lightly due to various environmental factors and inconsistent control over them. They mention that they wake up multiple times per night for different reasons like noise, temperature, mattress comfort, or their partner moving around. The author details several strategies they tried to improve their sleep quality, including stretching/yoga, mouth-taping, diet and hydration adjustments, sleep position changes, fan or white noise use, proper bedding materials, melatonin supplements, magnesium L-threonate, CBD oil, and weed. They also mention potential issues with some of these strategies, such as reliance on placebo effects, adverse side effects, or difficulty booking sleep studies due to travel restrictions.

Ig Nobel Prize Winner Higashiyama Atsuki and the "Between-Legs Effect" Mystery:

The Ig Nobel Prizes, often seen as a joke, have been awarded to various researchers for their work in unusual scientific fields. In 2016, Japanese researcher Higashiyama Atsuki received an Ig Nobel Prize for his study on the effect of viewing objects between one's legs (the "between-legs effect"). This

Last update: 2024/08/26 08:35

phenomenon involves bending at the waist and looking at one's surroundings upside down from a position between the legs. Higashiyama's research emphasizes the importance of considering physical information in psychology, challenging the idea that visual information alone is responsible for optical illusions.

Jhead:

Jhead is a command line driven program for manipulating the non-image parts of Exif flavour JPEG files that most digital cameras produce. It allows users to extract camera settings, adjust the time field, manipulate thumbnails, and edit metadata among other features. However, Jhead is limited in modifying certain pre-existing fields within the Exif header. The program also relies on other image processing tools such as ImageMagick for specific operations like generating thumbnails or cropping images.

Ekelhafter Influencer-Trend etabliert sich in Dubai:

The article discusses the alleged practice of influencers in Dubai being paid exorbitant amounts by wealthy men to participate in sex acts involving camels. This is known as the "Dubai Porta Potty" trend. According to reports, models are invited for lavish treatment and free flights to Dubai, where they are supposedly asked to engage in various depraved activities ranging from spanking to eating camel feces. One social media user claimed that a friend of hers was called to participate in a "Porta-Potty" party during her stay in the United Arab Emirates, in exchange for 50,000 Euros to engage in all possible fetishes. The details are alarming and can be read at the user's own risk.

Tech's broken promises: Streaming is now just as expensive and confusing as cable. Ubers cost as much as taxis. And the cloud is no longer cheap.:

The tech industry has been striving for profitability and raising prices to achieve it. As a result, various services are becoming increasingly similar to their predecessors. In video streaming, providers like Netflix, Disney, and Amazon are introducing bundles that cost as much as cable TV and even have ads in some instances. The original vision of streaming was cheaper and more convenient, but these changes indicate a shift away from those promises. Additionally, ride-hailing companies such as Uber and Lyft have increased prices while facing profitability goals. Lastly, the cloud's initial promise of cheaper and more secure computing for businesses is being challenged by rising costs and security issues.

Broken Ownership:

The "you build it, you own it" concept discusses the ownership trio of knowledge, mandate, and responsibility. It highlights that an individual cannot be responsible for something they don't control or understand. Broken ownership is more common than true ownership and leads to issues like frustration, demoralization, and poor operational performance. This article outlines six archetypes of broken ownership, including the monkey with a gun, foot soldier, baby parent, teenager, coma, and gambler scenarios. True ownership involves full control and understanding of the responsibility for the

2025/11/18 17:04 7/10 Links für 2023 KW 34

product or service.

A half-hour to learn Rust:

Thank you for the comprehensive explanation of Rust programming language features. It was very informative and helpful in understanding the syntax and features of Rust. Your writing style is easy to understand, which made the article enjoyable to read.

How to sabotage your salary negotiation efforts before you even start:

The article provides a comprehensive guide on how to approach salary negotiation during job interviews. It emphasizes the importance of preparation, having multiple offers before negotiating, and avoiding premature negotiation and revealing sensitive information to recruiters. The article also offers practical tips on handling recruiter calls, including remaining polite, expressing enthusiasm, and delaying the actual negotiation until all information is gathered. Overall, it provides valuable insights for individuals looking to effectively negotiate their job offer packages.

An Excruciatingly Detailed Guide To SSH (But Only The Things I Actually Find Useful) • Graham Helton:

This article provides a comprehensive guide to understanding and using various SSH flags for port forwarding, tunneling, and other advanced techniques. The author explains each flag, provides examples and illustrations of their use in real-world scenarios, and includes references for further reading on the topic.

It's Getting Too Hot for Tropical Trees to Photosynthesize, Scientists Warn:

The article discusses a study led by Christopher Doughty, which found that tropical leaves are approaching critical temperatures at which photosynthesis breaks down due to human-driven climate change. The researchers discovered that about 0.01% of leaves in the world's tropical forests may already surpass this critical temperature in a typical year, raising concerns about the fate of these essential ecosystems under pessimistic projections. If temperatures increase by approximately 3.9C, it could trigger a major photosynthetic meltdown for tropical forests, potentially causing plants to perish and leading to a collapse of this critical biome with far-reaching effects on climate and biodiversity. However, the authors stress that we can avert this disaster by rapidly decreasing our consumption of fossil fuels.

How to avoid KPI psychosis in your organization?:

KPI psychosis is a state where companies make decisions solely based on data and lose contact with reality. Companies are obsessed with collecting data to guide decisions, as it helps combat biases from human judgment. However, this can lead to tunnel vision, focusing only on what's measurable

while ignoring other important factors. To overcome KPI psychosis, companies should use a combination of intuition and data for optimal decision-making, ensuring that they consider the actual objective behind the metrics and regularly reflect on their reliability.

Close to the Machine: Technophilia and its discontents:

Ellen Ullman's memoir "Close to the Machine" delves into the programming process, detailing how it evolves from a clear-cut plan to managing chaos. The author emphasizes the immersive nature of programming and explores how it affects individuals both personally and professionally. Throughout the book, Ullman discusses topics such as revolutionary positions, trust issues, and the influence technology has on people's lives. "Close to the Machine" offers a unique perspective into the world of software development and computer culture, leaving readers with questions about the role technology plays in our lives.

Microsoft is bringing Python to Excel:

Microsoft is integrating Python into Excel to enhance data analysis and visualizations for users. A public preview of this feature is available now, allowing Excel users to access Python libraries directly from the ribbon in order to manipulate and analyze data within their spreadsheets. This integration aims to make advanced data analysis more accessible within the familiar Excel environment by leveraging Python's powerful capabilities directly from within the platform.

The Secret Weapon Hackers Can Use to Dox Nearly Anyone in America for \$15:

The article discusses how criminals are selling access to a powerful set of data, known as the target's credit header, which is personal information that Experian, Equifax, and TransUnion have on most adults in America via their credit cards. This information is valuable and often unprotected, allowing for its abuse by criminals. Criminals use Telegram bots to sell access to this data for \$15-\$40 depending on the type of information requested.

GitHub - chrieke/prettymapp: Create beautiful maps from OpenStreetMap data in a streamlit webapp:

The article discusses the "prettymapp" project on GitHub, which is a streamlit webapp allowing users to create visually appealing maps using OpenStreetMap data. It has received 1.2k stars and 56 forks and holds an MIT license.

How the Rich Get Richer:

The article discusses a study co-authored by economists from the International Monetary Fund (IMF) that finds wealthier individuals are more likely to earn higher returns on their investments. Data from 12 years of tax records in Norway reveals that high returns from wealth can prevent individuals from

2025/11/18 17:04 9/10 Links für 2023 KW 34

leaving the wealth distribution's top percentages. Wealthy individuals may earn higher returns, even on more conservative investments, due to access to exclusive opportunities or better wealth managers. The study also shows that while wealth can be passed down to successive generations, high investment returns may not persist across generations.

Shrinking Economies Don't Innovate:

The article discusses how innovation plays a crucial role in economic growth, with competing factors influencing its rate. As population declines, innovation slows down, resulting in a shrinking economy that struggles to restart growth using tech or innovation solutions. This could lead to a decrease in world productivity and mood levels, increasing conflict and environmental impacts. To maintain growth, the article suggests that capital, output, and tech levels must double every generation. However, current global fertility rates raise concerns about reaching this doubling threshold.

Latest Android Runtime (ART) update led to apps starting 30% faster:

The Android Runtime (ART) is an essential component of the Android operating system that provides runtime and core APIs for apps and most OS services. It has been updated via Google Play system updates since Android 12, offering faster app startup times, execution speed improvements, enhanced memory usage, improved bytecode compilation, and security fixes. The latest ART update, version 13, provided "improvements of up to 30% on some devices." With future updates like ART 14 rolling out in the coming months, Google will continue optimizing performance while reducing code size. These updates are available on Android 12 and newer devices, with plans for Android Go soon.

Texas electricity prices soar 6,000 percent:

On August 18, 2023, Texas power prices increased by 6,000%, reaching \$4,750 per megawatt-hour due to a new heatwave. This rise in temperature will likely break records for the area. Spot electricity prices nearly reached the cap of \$5,000. The high temperatures are expected to persist over the weekend with values ranging between 105 and 113 degrees. Texas power market is deregulated, and its grid has faced massive spikes in demand this summer, but supply remains steady due partly to renewable energy sources such as wind and solar.

De Facto Ports:

This article discusses the use of port numbers in applications and services. It mentions that ports above 1024 are generally available for use, while ports below 1024 often require special privileges. It presents a list of commonly used port numbers and their associated protocols, noting that even numbered ports tend to be reserved for development purposes, while odd numbered ones are more likely for production use. The article also observes that there seems to be no central theme or dominant application using any particular port number.

Last update: 2024/08/26 08:35

GitHub - ozmartian/vidcutter: A modern yet simple multi-platform video cutter and joiner.:

The article discusses vidcutter, a multi-platform video cutter and joiner with a GPL-3.0 license. It has received 1.6k stars and has been forked 131 times on GitHub.

Anxious individuals shift emotion control from lateral frontal pole to dorsolateral prefrontal cortex - Nature Communications:

This article discusses a study that investigates the differences in the neural circuitry responsible for controlling emotional action tendencies between high-anxiety individuals and their non-anxious counterparts. The study utilizes neurochemical, structural, and functional measures to characterize neural circuits supporting emotional control. The findings indicate that anxious individuals tend to use dIPFC rather than FPI in implementing control over emotional action tendencies, and that FPI in high-anxiety individuals may receive stronger input from the amygdala via more extensive amygdalofugal pathway connections. Additionally, the study reveals that FPI in high-anxiety individuals is highly excitable and that the excitation/inhibition balance of that region is decoupled from behavioral and neural indices of emotional action control. These findings offer a mechanism for explaining how emotional challenges in anxiety disorders can generate a neural bottleneck in the control of emotional action tendencies.

Physics — Susan Rigetti:

Thank you for providing this summary! It seems to be a comprehensive guide to learning physics on one's own, with a wide variety of textbooks and supplementary resources suggested throughout. This is valuable information for anyone interested in self-studying physics, especially for those who might not have access to formal educational opportunities.

links, 2023

From:

https://tobias-fink.net/ - Tobis Homepage

Permanent link:

https://tobias-fink.net/content/2023/links-kw34

Last update: 2024/08/26 08:35

