Foundation, Frontier, or Large Language Models: Why names matter and a case for embracing the tradition of "Muppet Models"

Anonymous ACL submission

Abstract

One of the trends of AI is the creation of machine learning models that can fluently generate coherent, useful language. But what should we call these models? LLMs, Foundation Models, Stochastic Parrots, Frontier Models are possible options. . . but years ago we had a consensus: many of the progenitors of today's superstar AI models were named after Muppets. This trend started as an inside joke and was later used as a way to pay homage and build on each other's work, but has collapsed as AI has become mainstream. Analyzing over 150 models, we identify patterns across Western and Chinese traditions, showing how names reflect values, branding, and trust. We argue that naming is not trivial but central to how AI systems are framed. And we make the case for reviving the playful metaphor—balancing hope and skepticism—of the Muppet Model (MM).

1 The Name of the Model

In an accident of history, the first large language models were named after Muppets, as we review in Section 3. However, the community has mostly moved on from this tradition (Section 4), and different traditions have emerged for naming these models in the West and in China: today's names come from gods, generic human names, sci-fi characters, and technical descriptions.

This position paper argues that the perception of AI is influenced by what models are called (Section 2), and that we as a field should both be aware of this influence and be responsible researchers by naming general and specific models wisely.

As the tragic story of teenage suicides after forming relationships with chatbots show (Chatterjee, 2025), human perceptions of models have consequences. Beyond documenting what these models are named, the contribution of this position paper is to argue that the original naming of "Muppet" models was apt (Section 6): calling large, generative AI systems Muppet Models is a good compromise for a new technology with limited skeuomorphism and better than the alternatives (Section 5).

We are not so fanciful as to expect the Techcrunch will create a "Muppet Moves" beat to track AI news. Instead, the goal of this position paper is to have a hot enough take that will help the reader: learn (or refresh) the history of these models through the lens of their names; discover patterns in how models are named today (and how China is different); question the hype-infused naming conventions that dominate the field today; think critically about the increasing role these models play in society; and reflect that maybe we got it right in the first instance by calling them Muppet models.

2 The Name's the Thing

Gardner and Levy (1999, p. 131) contend a

> name is more than the label employed to differentiate among the manufacturers of a product. It is a complex symbol that represents a variety of ideas and attributes. It tells the consumers many things, not only by the way it sounds (and its literal meaning if it has one) but, more important, via the body of associations it has built up.

This is more important when products do not distinguish themselves– as is the case with the convergence of AI models. With the reader's indulgence the rest of the paper will stand by our thesis and call these models **Muppet models** (MM for short).

2.1 The Power of a Name

In his tract "The Power of a Name", Graham (2013) argues that naming a scientific concept is an awesome power: by merely providing the name to multiple infinities, Cantor (Zipes, 2015) changed the field's understanding as much as with his proofs, an extension of the fables where knowing the name of a spirit offered control over them (e.g., Rumpelstiltskin). Mathematician Alexandre Grothendieck (2021, p. 90) equated mathematical research with naming, imagining

> unknown and nameless worlds, crying out for me to become acquainted with them and bestow names upon them.

Moreover, a brand name is essential for building brand recognition, protecting business interests, and influencing consumer psychology and perception (Kottler and Keller, 2009; Keller, 2013). An appropriate name and branding strategy help a company stand out in the market and maintain long-term competitiveness and sustainability (Bulatovic et al., 2016). If users like GPT-3, they will trust GPT-4 and expect it to be better. This continuity builds trust and anticipation among users.

3 In the Beginning were Muppets

This section discusses the first MMs which were actually named after Muppets, explains *why* they were named after Muppets, and then ends with the close call by OpenAI to *not* name GPT after a Muppet and how that lead to the explosion of non-Muppet names (Section 4).

3.1 The Muppet Age: ELMo, BERT, and the Unseen Snuffy

The Muppets are puppets developed by Jim Henson for a the Washington, DC show *Sam and Friends* (Jarvis, 2015). Unlike their progenitors, Muppets interact directly with humans on scripted children's education television like *Sesame Street* (Gettas, 1990), variety shows for adults like *The Muppet Show*, and in feature films like *The Muppet Movie*. Because of their ubiquity, Muppets provide names that are simultaneously familiar, memorable, and lightly humorous.

To the best of our knowledge, the first model to intentionally adopt a Muppet name used embeddings from language model (Peters et al., 2018b, ELMo), a bidirectional LSTM whose representations could be adopted for downstream tasks.[1] Unlike the previous seismic shift in NLP—WORD2VEC (Mikolov et al., 2013)—ELMo created distributed representations of tokens *in context*. Shortly thereafter, bidirectional encoder representations from transformers (Devlin et al., 2019a, BERT) used the newly released transformer architecture (Vaswani, 2017) to combine most of the ingredients we associate with modern models.

In addition to these "canonical" MMs, similar models emerged, reinforcing the trend: ERNIE (Zhang et al., 2019a), KERMIT (Zanzotto et al., 2020), BIGBIRD (Zaheer et al., 2020b), ROSITA (Mulcaire et al., 2019), and GROVER (Zellers et al., 2019a).

3.2 The Great Break: From Snuffy to GPT

The break from the dominance of using Muppet names came when OpenAI released the generative pre-trained transformer (Radford et al., 2019, GPT2), which originally had a Muppet name: "Snuffleupagus" ("Snuffy" for short). Clark (2019, (posted on X)) believed it

> poor taste to give muppet [sic] name while discussing reasons to be cautious with regard to increasingly powerful language models

Thus, GPT emerged as a more sober, descriptive brand– marking the start of the era of neutral, non-Muppet names that dominate today. The choice marked a turning point.[2]

Later Western models followed this more technical naming trend. Most notably, many new models adopted GPT-derivative naming conventions, such as domain- or capability-specific variants (e.g., BioGPT, FinGPT, CodeGPT). In our dataset, 30% (N=157) of post-2019 Western models used GPT-derived or GPT-inspired naming patterns, illustrating how the success of the GPT family influenced subsequent naming practices. In contrast, a panoply of new Chinese models offered other competing naming schemes, as we outline in the next section.

4 What are models called today?

Since the release of GPT-2 in 2019, model names have increased their diversity. After outlining our inclusion criteria, this section examines what models are called today (as of 2026), why they bear those names, and what broader cultural, institutional, and market forces these naming conventions reflect.

4.1 Selection Criteria for Muppet Models

While we will define what we specifically mean by Muppet models as part of our definition of the models in Section 6, for now we outline a more generous definition that encompasses what the field generally considers LLMs, frontier models, etc. (We

[1] Given the definitions in Section 6, it's arguable that ELMo does not meet our definition today. However, it deserves honorary status for historical reasons. One could also argue that some models like DAN (Iyyer et al., 2015) were accidentally named after Muppets.

[2] At the risk of editorializing, Snuffleupagus (a giant Muppet that only a select few can see) is a good name for a model released by an increasingly secretive "Open"AI pursuing hyperscaling.

discuss these alternative umbrella naming schemes in Section 5.) To be included, a model must either (a) generate coherent and contextually appropriate text, or (b) form a distributed representation of compositional meaning that enables text generation or downstream linguistic reasoning.[3] Moreover, to be included a model must satisfy at least two of the following "notoriety" standards:

- **Published paper or technical report.** Models must have publicly accessible technical documentation describing their architecture, training approach, capabilities, or deployment in sufficient detail to support scholarly analysis and attribution. Examples include Baichuan 2 (Bai et al., 2023), LLAMA (Touvron et al., 2023a), and Gemini-1 (Team et al., 2023), which provide formal reports describing model design, training, evaluation, and intended use.
- **Public availability.** Models must be released through recognized repositories or institutional platforms with verifiable provenance and sufficient public access to support inspection, attribution, and reproducibility. This criterion emphasizes accessibility of model artifacts (e.g., weights, tokenizers, technical documentation, or public API access) rather than a strict open-versus-closed distinction. Examples include DEEPSEEK-V3 (DeepSeek, 2025), MISTRAL-7B (Vonplaten, 2025), and Grok (xAI, 2025), which provide publicly accessible model artifacts or APIs alongside technical descriptions. Conversely, proprietary or domain-specific GPT variants were excluded when their artifacts or APIs were not publicly accessible or independently inspectable.
- **Large user base or industry deployment.** A model can qualify if verifiable evidence shows large-scale adoption or integration into widely used platforms. GROK has 30.1 million monthly active users and 153 million website visits (Cardillo, 2025), thus qualifying under public availability and large user base.

We require two of these three conditions to remove unverifiable and fringe models. Appendix Table 1 presents the complete census of models, including those that did not meet the inclusion threshold, For example, internal GPT-4 training checkpoints were excluded because they lacked public accessibility and verifiable documentation, while models such as StructBERT were excluded despite having technical documentation because they lacked broad public deployment or visibility. Appendix Table 2 summarizes qualifying models' inclusion criteria.

4.2 The Rainbow Collection

Language model names fall into several general categories, each reflecting different priorities and cultural resonances. Section 3 introduced **Muppet-inspired** and **Capability-driven names**, and Muppets are an extension of a broader trend of **Popular culture references**: BART (Lewis et al., 2019), Megatron (Shoeybi et al., 2019). Because science is as much evoluation as innovation, **Derivations from original** build on existing names: (Liu et al., 2019a, ROBERTA); (Khattab and Zaharia, 2020, COLBERT). Some models adopt **personal, historical, or scientific names** to evoke familiarity, intelligence, or cultural prestige: Claude (Anthropic, 2025) and Einstein (Salesforce, 2025). By contrast, CAIMIRA (Gor et al., 2024) draws on **mythological references**. **Geographical names** draw on locations or landmarks (Sengupta et al., 2023, Jais), and **Company branding** foregrounds institutional identity (Bi et al., 2024, DeepSeek). **Flora and fauna**, combining natural imagery with technical or metaphorical meaning (Touvron et al., 2023a, LLAMA). Together, these categories illuminate the rich strategies researchers and developers use to name models, balancing memorability, technical signaling, and cultural resonance (Appendix Section B).

4.3 Annotator Agreement

Two authors[4] independently assign these categories to 157 models, which were then validated by two additional independent annotators using the same predefined category definitions. Agreement is high: 150 out of 157 matches (95.5%), yielding Cohen's $\kappa \approx 0.92$, which qualifies as "almost perfect agreement" (Landis and Koch, 1977). Further details

[3] This functional criterion allows inclusion of models such as BERT and ELMO, which do not generate text directly but are foundational to generative architectures, while excluding narrowly task-specific systems such as classifiers or regressors that sometimes are termed "AI" in the popular discourse.

[4] We did not use crowdworkers, since the task requires interpretive judgments about naming origins, familiarity with model naming conventions, and consistent application of a nuanced taxonomy, which are difficult to ensure in a purely crowd-sourced setting.

are in Appendix D. Disagreements were rare but revealing. For instance, GEMMA was classified by one annotator as an English personal name (reflecting a human or Western given name), while the other labeled it as a conceptual name, interpreting it as derived from the Latin word gemma, meaning "precious stone" or "gem".

4.4 Trends in Western Model Naming

Western LLM naming has evolved in three overlapping phases, reflecting research culture and commercialization priorities. Between 2018–2020, playful and pop-culture-inspired Muppet names dominated, accounting for approximately 70% of models in this period (but only 22% overall). This phase includes influential Muppet models such as (Devlin et al., 2019b, BERT) and (Peters et al., 2018b, ELMo).

From 2020–2022 there was an inflection. Models were not created *de novo*: they built on existing systems. However, since prior models already established this style, playful naming persisted, with names often combined with predecessors to clarify function, representing roughly 44% of Western models. Muppet derivatives include DistilBERT (Sanh et al., 2019), albert (Lan, 2019), and DeBERTa (He et al., 2021). At the same time, playful names like FLAN (Wei et al., 2022) and PaLM-E (Driess et al., 2023) remained, though more models adopted direct technical names such as T5 (Xue et al., 2021) and UL2 (Tay et al., 2023). These names emphasize lineage, architectural improvements, and specialization, enabling clearer tracing of technical progress.

After 2022, commercialization and global branding drove a rise in scientific, conceptual, and nature-inspired names (≈22%). Names such as Jurassic (Lieber et al., 2021), Galactica (Taylor et al., 2022), Mistral (Jiang et al., 2023a), Nova (Jiang et al., 2025), Zephyr (Tunstall et al., 2024), Chinchilla (Hoffmann et al., 2022), Sparrow (Glaese et al., 2022), and Falcon (Almazrouei et al., 2023) are less "playful" but more formal, conveying scale, sophistication, and metaphorical associations with performance.

4.5 Trends in Chinese Model Naming

Chinese LLM naming trends are similar in that they have become more serious over time. With a few prominent exceptions, Chinese models shifted away from playful or mascot-inspired naming toward mythology and traditional cultural references.

Historical and mythological references account for approximately 44% of Chinese models in our dataset released between 2021–2024: Pángǔ (Zeng et al., 2021, a mythical figure who separated heaven and earth), Xuānyuán (Zhang and Yang, 2023, the birth name of the mythical "yellow" emperor), and Mèngzǐ (Zhang et al., 2021b, aka Confucius). These names embed models within culturally recognized narratives, signaling heritage, symbolic depth, and alignment with classical literature and philosophy (Wu, 2016; Yu, 2023).

Earlier Chinese models were more frequently capability-driven or functional in naming (e.g., CPM, EVA, ChatGLM), but between 2022–2024 naming expanded toward traditional concepts (approximately 24% of Chinese models released during this period) and geographical or place-based identifiers (approximately 9% during this period). Examples include Bǎichuān (Yang et al., 2023, a bigram from a proverb meaning "all rivers flow to the sea"), Wù Dào (BAAI, 2021, the path to enlightenment), and Hùnyuán (Sun et al., 2024, the fundamental and unified source from which all phenomena and energies arise in Daoism). This reflects both aesthetic/philosophical considerations and regulatory legacies, such as requirements for geographical indicators in enterprise names (SAIC, 2012). From 2024–2025, increasing international attention encouraged more scientific and conceptual naming conventions, representing approximately 24% of Chinese models released during this period: DeepSeek (Bi et al., 2024), SEEChat (LLC, 2025), and TeleChat (Wang et al., 2024c).

4.6 How Naming Patterns Differ by Region: Statistical Analysis

The differences between Chinese and Western LLM naming reflect underlying sociocultural and regulatory influences. In the West, playful Muppet-style names ($\approx$ 22% of Western models in our dataset) evolved into lineage-focused technical descriptors ($\approx$ 44%), then globally resonant conceptual names ($\approx$ 22%). Recent media discussions note the complexity of naming conventions, with OpenAI CEO Sam Altman acknowledging that labels such as GPT-4.1, "mini," or "nano" can be confusing (Perkel, 2025).

In China, naming reflects explicit government influence, classical educational traditions (Wu, 2016; Yu, 2023), and political discourse promoting cultural confidence (Xi Jinping, 2025). Mythological, historical, or literary names dominate early mod-

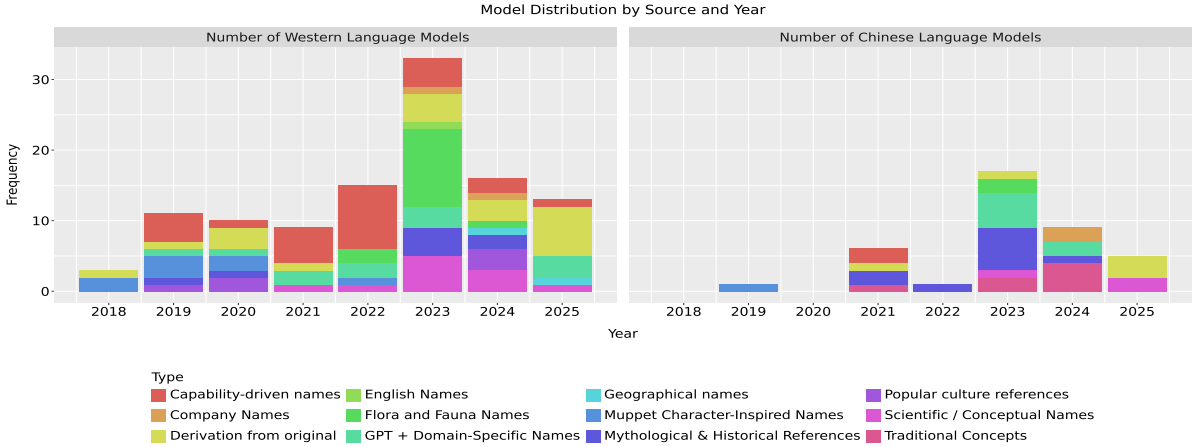

Figure 1: **Evolution of LLM Naming Patterns in China and the West (2018–2025).** Western models emphasize acronyms, derivations, and technical descriptors, while Chinese models favor cultural and symbolic references.

els, complemented by geographic or traditional concept-based names in later phases. Some models straddle this divide by having different Chinese and English names. Baidu's ERNIE retains a Muppet-style English name, while its Chinese name ("Wénxīn Yīyán") adopts a more poetic domestic framing. Creators systematically leverage cultural references to differentiate models (Jiang et al., 2023b; Lemos et al., 2024), aligning with broader Chinese brand-naming conventions (Chan and Huang, 2001).

Chi-square analysis confirms significant regional differences in naming types ($\chi^2 = 34.9$, $df = 10$, $p < 0.001$, Cramer's $V = 0.59$). Category-level contrasts further highlight these differences: traditional figures/concepts comprise 58.1% of Chinese models compared with 25.0% of Western models; GPT + domain variants account for 19.4% of Chinese models vs 10.3% of Western models; Muppet-style acronyms 3.2% vs 11.8%, respectively; and geographical or place-based names 3.2% for Chinese models vs 0% for Western models (Figure 1). Western models with surface-level geographical connections- such as LLAMA, FALCON (animal) and MISTRAL (wind)- were coded under fauna and traditional concept categories, respectively, following stricter classification criteria, even though their names retain some indirect geographical associations. Naming is both functional and culturally legible: Western models increasingly prioritize global marketability and technical clarity, while Chinese models embed symbolic, historical, and cultural legitimacy alongside emerging international branding strategies (Schroeder et al., 2015). As Bowker and Star (2000) observe, classification systems reveal "deep structures of power and practice".

5 Why is our Frontier not a Foundation?

This section evaluates three widely used alternatives: foundation models, frontier models, and stochastic parrots. None fully captures the balance between capability, flexibility, and risk.

5.1 Foundation Models: Presumed Stability

The term *foundation models*, popularized by Bommasani et al. (2021), evokes a structural metaphor: these models serve as a base upon which future applications and specialized systems can be built. The image implies stability— you *should* build on these models— but also adaptability: by fine-tuning (Lu et al., 2025), adding modalities (Sampath et al., 2025), or modifying objectives (Wu et al., 2025), new models can emerge while retaining shared representations. Yet this framing, while elegant, also risks overstating solidity: you trust a foundation but you usually take it for granted. . . it will always be there. As we outline in the next section, we believe our collective name for these models should emphasize interaction and fluidity.

5.2 Frontier Models: Vague, Undeniable Techno-optimism

The term *frontier models* has gained traction in policy and safety discourse to describe cutting-edge AI systems and their associated risks. Yet, as Bullock et al. (2024) note, there is no consistent definition of frontier models, and the term is sometimes deployed more for rhetorical effect than analytic clarity. Unlike *foundation models*, which emphasize stability and reuse, *frontier models* foreground novelty, exploration, and uncertainty—often at the cost of conceptual precision or regulatory accountabil-

ity. The metaphor carries deep cultural resonance: Turner's "American frontier" evoked opportunity and expansion beyond known boundaries (Turner, 1893); Kennedy's "New Frontier" tied the term to scientific ambition (Kennedy, 2025); and *Star Trek* cast space as "the final frontier," (Pearson et al., 2014) a site of wonder and discovery. Yet, as McLure (2000) warns in *The Wild, Wild Web*,

> because the electronic frontier is still generally a lawless territory, vigilantism is often the preferred- and sometimes the only effective- response to what the cybersettlers perceive as crimes against both property and people.

Applied to AI, this metaphor thus captures both promise and peril: frontier models evoke adventure and innovation but also imply danger, opacity, and a lack of social constraint.

5.3 Stochastic Parrots: Accurate but Overly Dismissive

Stochastic parrots, introduced by Bender et al. (2021), underscore language models' dependence on statistical correlations rather than genuine understanding. This metaphor has become central to responsible AI debates- raising concerns about the environmental costs of large-scale model training (Strubell et al., 2020; Ganguli et al., 2023), the data exploitation and social biases embedded in training corpora (Abid et al., 2021), and the ethical risks of anthropomorphization and misplaced agency (Weidinger et al., 2021). It remains a crucial corrective to hype-driven narratives that overstate machine "intelligence". At the same time, this framing risks flattening the complexity and creative affordances of Muppet Models (MMs). As Mitchell (2025) observes, models "can be useful without being trustworthy". Like Bender et al. (2021), we caution against taking MMs too seriously: they are synthetic, constrained, and should not be mistaken for sentient or moral agents. Yet we also recognize that they are fun, expressive, and often genuinely useful. Our name should reflect both truths.

6 Reject Modernity, Embrace Tradition

Having discussed what models are currently named, we turn from a descriptive accounting to MMs to justify why the models should still be called MMs.

6.1 What is a Muppet Model

A Muppet Model is a type of large, generative AI system built around language but not limited to it, that embodies several core traits.

Multimodal, Language-Centric Design. Muppets talk,[5] and MMs are fundamentally language-focused, but they can go beyond just text. Some, like VISUALBERT (Li et al., 2019) and VILBERT (Lu et al., 2019), combine language with images, allowing them to understand and produce multimodal content.

6.1.1 Controllability and Alignment

The central conceit of a Muppet is that it is a character. You are not talking to a piece of felt with ping pong balls glued to it. Rather, you are talking to Kermit, who has a history and a personality. That is an artifice that has been built up, and the felt and ping pong balls contribute to that. But there is a man behind the Muppet! Muppets are controlled by a hand inside their head, guiding their actions and speech. This metaphor captures central challenge of alignment in large models: ensuring outputs reflect human goals, constraints, and values rather than drifting into unguided or harmful behavior (Amodei et al., 2016; Leike et al., 2018; Bender et al., 2021). MMs are not passive generators of text; they can be actively steered. Techniques such as prompt engineering (Ouyang et al., 2022; Leike et al., 2022), task-specific fine-tuning (Ziegler et al., 2019), and reinforcement learning from human feedback (Christiano et al., 2017, RLHF) serve as the "hand inside the Muppet", constraining otherwise open-ended generation.

Just as a Muppet never speaks without a puppeteer, MM do not act independently: their actions can be shaped to follow a constitution, rebel against woke strictures, or avoid "off limits" topics. Anthropic's Constitutional AI framework, for example, uses explicit high-level principles to guide model responses and critique unsafe behavior (Bai et al., 2022). Other systems are marketed around alternative alignment philosophies, such as xAI's Grok, which was explicitly framed as less constrained by "woke" moderation norms (Waight et al., 2026; Fox News, 2026). Meanwhile, models developed under stricter domestic regulatory regimes often avoid politically sensitive topics or decline to answer them entirely, reflecting either alignment objectives or compliance constraints (Noels et al., 2025). Emphasizing the human agency with the Muppet analog makes this clearer.

[5] Even if sometimes it's hard to understand or gibberish disguised as Swedish

Yes, And Muppets. Muppets are not static puppets reciting a fixed script; they are increasingly deployed in the real world. Jim Henson emphasized improvisation as central, noting that Muppets succeed when they feel alive—reacting in the moment rather than merely delivering prewritten lines (Guimarães, 1997). Improvisation and interactivity define Muppet performances. Kermit the Frog frequently "breaks the fourth wall", engaging viewers or responding to unpredictable events (muppetshow, 1976).This is consistent with MMs, for better or for worse. The models sometimes react to human suggestions with sycophancy (Wei et al., 2023; Sharma et al., 2023), but also help build alluring worlds in creative settings (Chung and Kreminski, 2024).

6.2 Why it Remains a Good Name

Having traced the shift away from MMs as a naming convention, we acknowledge the momentum of historical and market trends. Nevertheless, we argue that there remains a compelling case for preserving the tradition of referring to these increasingly ubiquitous AI models as MMs.

Muppets are fun, but not too serious Crawford (2021) makes the case that framing AI systems through names that invoke playfulness or metaphor, rather than exaggerated competence, can temper user overreliance and prompt critical engagement.

Unlike other names for models that focus on their advanced features or importance, "Muppet" evokes fun more than competence. Indeed, in the UK, "muppet" often refers to someone behaving foolishly, which likewise undercuts assumptions of competence (Crystal, 2000). In our opinion, this accident of history–a neutral association with unpredictable characters designed for entertainment–is a more accurate characterization of models' capabilities than boosterism (Frontier models) or minimizing names (stochastic parrots).

Just as viewers of *Sesame Street* accept most of what they hear while remaining aware that it is a performance, the name MM does not invite unquestioning trust. Research in cognitive framing and human–computer interaction shows that labels and metaphors shape user trust (Lee and See, 2004; Rathi et al., 2025); in this light, the playful "Muppet" metaphor signals that these systems are powerful but limited, helping temper user overreliance.

While large language models have profound societal impacts and raise significant safety, ethical, and governance concerns, our argument is not that these systems are trivial. Rather, a playful metaphor serves as a counterbalance to names that imply excessive competence or authority. "Muppet" may remind users that these systems remain constructed artifacts rather than dependable people. Importantly, we do not propose replacing technical specification with whimsical naming; rather, evocative public-facing names can coexist with precise architectural descriptions in papers and documentation.

A Technical Tour-de-Force. Muppets, like modern models, are a combination of many technical innovations: from the armature enabling expression (Kaestle et al., 1980) to underwater puppeteers (Finch, 1993), it takes a lot to make a Muppet function. Likewise, the innovations in scale and algorithms in MMs are also remarkable. Using MM helps retain the wonder of invention without forgetting that these models are an artifice (Huang et al., 2025; Du et al., 2022a).

Multicultural and Multilingual Dynamics. Although Muppets originated in English, they have been adapted into many languages and cultures. Like all cultural metaphors, their meaning depends on shared context and may not be interpreted identically across audiences. *Sesame Street* has been localized into over 30 international versions (Gettas, 1990; Kerley, 2013), each introducing characters tailored to local cultures. The German adaptation, *Sesamstraße*, features a large brown bear absent from the American cast, illustrating how Muppets adapt to local contexts (Fandom, 2025). As the producers emphasized, they do not simply "translate" the program but transform it for each setting. That said, Muppets are still American. A case in point is *Fraggle Rock* (Henson, 1983)–a show with multiple literal layers with Humans, Doozers, Gorgs, and the titular Fraggles. The show used local actors for the French, UK, and German introductions. This slightly reframed the "inner" (below ground) story with the actual Fraggles, while the main story remained a largely literal translation of the American version. This provides a useful metaphor for contemporary language models. While these systems can interact across languages and cultures, meaningful adaptation remains non-trivial and largely limited to well-resourced languages and widely represented cultures (Joshi et al., 2020; Bender et al., 2021). As with the layered structure of *Fraggle Rock*, only surface elements may be adjusted

for local audiences, while much of the underlying narrative remains unchanged. Modern models are similarly only "tailored" in limited ways: prompting, deployment context, or interface design may adapt some interactions, but most behavior still reflects culturally dominant patterns present in training data.

Learning From Fantastical Creatures. Muppets can also be educational tools. Like how RAG (Lewis et al., 2020c) can make models consistent with underlying text, the Muppets have a tradition of staying close to the source text, as Napolitano (2009, p. 79-102) notes, the *The Muppet Christmas Carol* strives

> to preserve key elements of the original text, thus reinforcing the appeal of Dickens's story and ultimately encouraging young viewers to seek out the source itself.

But despite often very high fidelity to the source and their desire to be educational, Muppets, by their nature, cannot always be trusted. The movie also contains the line (Henson, 1992)

> Know it? My first job was here. This is Fozziwig's old rubber chicken factory,

which our research suggests was not part of Dickens's original story.

Sigler (2014) noted how Muppets— in particular the Fraggle Rock character "Der Allwissende Müllhalde" (the All-Knowing Trash Heap)— had become synonymous (even before the AI) in German online slang with misinformation. An oracle whose answers were often wrong yet delivered with conviction remains an apt metaphor for search results from unreliable online sources. Modern language models share this paradox: contextually fluent, educational, and persuasive, yet fallible (Example in Appendix Figure 2).

Corporate Overlords. Perhaps our most persuasive reason for using the term Muppet models is also the most practical reason not to: in 2004, Disney acquired rights for the "Muppets" properties from Henson (Disney Press Relations, 2004). Disney is a large, multinational corporation. Like the purveyors of MMs generally, they have their own agendas and can make the people who want to use the the term "Muppets" without their explicit permission have a bad time. Likewise, the people who put out MMs have their own agendas and political goals, which we should not forget when we ask Grok about the results of an election (Hajj, 2024). While we think this is orthogonal to the academic argument of this position paper, we do present a workaround in Appendix A.

Limited Skeuomorphism and the Case for "Muppet" MM's limited skeuomorphism (when knowledge of one device translates to another) has been a challenge: what is this thing you're talking to? A teacher? A search engine? A friend? A priest? A servant? Unlike the other potential names, we have no guidebook about how to talk to a Frontier or a Foundation. And while you might talk to a Parrot, it won't be a productive conversation. While many people have never talked to a Muppet, this makes the name a more accurate reflection of the model's capabilities: interactive, improvisational, and guided, yet not foundational or lawless. It will be fun, probably useful, but something you should double check.

7 Conclusion

This paper traces the evolution of the Muppet Model naming tradition– from playful beginnings to a lasting influence on how large language models are conceived, discussed, and trusted. Despite different trajectories, both Chinese and Western names for models have converged to sterile, non-objectionable monikers. That is because AI today is serious business. But that doesn't mean that we should take MMs or their outputs too seriously. Indeed, today there is even more reason to embrace older traditions. These new models are useful, fun, exciting, and towering achievements of technical brilliance, but nonetheless an artifice that— no matter how much they sound like a person—- somebody's hand is guiding what's coming out. And behind that puppeteer is often a big corporation with its own agenda.

As scientists, we have a duty to both discover new concepts but also to be good shepherds of our ideas. Responsible names are a part of that duty. MMs remind us that these systems are built to converse, interpret, and collaborate—not merely compute. Preserving this lineage maintains a vocabulary that is both technically meaningful and socially intelligible, bridging algorithmic complexity and public understanding. And it can help make sure that we also don't take ourselves too seriously, preserving the spirit of the innovators who helped usher in the era of ELMo and BERT.

8 Limitations

Our analysis is necessarily partial and interpretive. The landscape of MMs evolves too rapidly for any census to remain fully comprehensive;

new systems are released, renamed, or rebranded continuously. Despite extensive cross-referencing with open databases and corporate announcements, some models may have been missed or misclassified.

Second, while the Muppet metaphor offers a productive lens for understanding mediated agency and co-performance, it remains culturally situated. Muppets originated in the United States and carry Western assumptions about creativity, humor, and personhood. Although the metaphor's flexibility allows cross-cultural adaptation, it may not fully capture naming or design logics grounded in non-Western traditions—especially those emphasizing collective identity, cosmological symbolism, or institutional authority.

Third, our analysis relies on official or publicly documented names. We do not include informal, user-assigned, or emergent names that models acquire in deployment contexts (e.g., community nicknames or internal code names). These unofficial identifiers may reveal additional cultural and affective dynamics that fall outside formal naming conventions.

Another limitation of this work is that the categorization of model names relies on annotator interpretation rather than large-scale user studies. Although inter-annotator agreement was high and additional annotators showed substantial agreement with the final inferred associations, these annotations may not fully capture how broader populations perceive or interpret model names across different linguistic and cultural contexts. Ideally, future work would involve surveying a larger and more diverse set of participants to evaluate whether the associations identified in this paper align with public perception. Such studies could provide a more comprehensive understanding of how naming conventions influence user expectations and interpretations of AI systems.

References

Abdulhady Abas Abdullah, Arkaitz Zubiaga, Seyedali Mirjalili, Amir H. Gandomi, Fatemeh Daneshfar, Mohammad Sadra Amini, Alan Salam Mohammed, and Hadi Veisi. 2025. Evolution of Meta's LLAMA models and parameter-efficient fine-tuning of large language models: a survey. *CoRR*, abs/2510.12178.

Abubakar Abid, Maheen Farooqi, and James Zou. 2021. Persistent anti-muslim bias in large language models. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 298–306.

Aaron Adcock, Aayushi Srivastava, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pande, Abhinav Pandey, Abhinav Sharma, Abhishek Kadian, Abhishek Kumawat, Adam Kelsey, and 1 others. 2026. The llama 4 herd: Architecture, training, evaluation, and deployment notes. *arXiv preprint arXiv:2601.11659*.

Alibaba. 2025. Cao Zhi Large Model-Focus on long-form text, multilingualism, and verticalization — aibase.com. https://www.aibase.com/tool/28994. [Accessed 14-02-2025].

Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Mérouane Debbah, Étienne Goffinet, Daniel Hesslow, Julien Launay, Quentin Malartic, and 1 others. 2023. The falcon series of open language models. *CoRR*.

Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. 2016. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*.

Anthropic. 2025. Anthropic's Transparency Hub: Model Report. *Anthropic*. Online; Available at: https://www.anthropic.com/transparency/model-report.

Kirolos Ataallah, Xiaoqian Shen, Eslam Abdelrahman, Essam Sleiman, Deyao Zhu, Jian Ding, and Mohamed Elhoseiny. 2024. Minigpt4-video: Advancing multimodal llms for video understanding with interleaved visual-textual tokens. *CoRR*, abs/2404.03413.

BAAI. 2021. Beijing academy of artificial intelligence-wudao. *GitHub*. Online; Available at: https://github.com/BAAI-WuDao.

Mingliang Bai, Zhihao Zhou, Ruidong Wang, Yusheng Yang, Zizhen Qin, Yunxiao Chen, Chunjin Mu, Jinfu Liu, and Daren Yu. 2023. Houyi: An open-source large language model specially designed for renewable energy and carbon neutrality field. *Preprint*, arXiv:2308.01414.

Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen, Yanru Chen, Yuankun Chen, Yutian Chen, Zhuofu Chen, Jialei Cui, Hao Ding, Mengnan Dong, Angang Du, and Chenzhuang Du. 2025. Kimi K2: open agentic intelligence. *CoRR*, abs/2507.20534.

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, and 1 others. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.

Baidu. 2025. Introducing PCL-BAIDU (Wenxin ERNIE 3.0 Titan), the World's First Knowledge Enhanced Multi-Hundred-Billion Model. *Baidu Research*. Online; Available at: https://www.reuters.com/world/china/alibaba-launches-qwen3-max-ai-model-with-more-than-trillion-parameters-2025-09-24/.

Marco Bellagente, Jonathan Tow, Dakota Mahan, Duy Phung, Maksym Zhuravinskyi, Reshinth Adithyan, James Baicoianu, Ben Brooks, Nathan Cooper, Ashish Datta, and 1 others. 2024. Stable lm 2 1.6 b technical report. *arXiv preprint arXiv:2402.17834*.

Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623.

Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiushi Du, Zhe Fu, and 1 others. 2024. Deepseek llm: Scaling open-source language models with longtermism. *arXiv preprint arXiv:2401.02954*.

Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O'Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, and 1 others. 2023. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning*, pages 2397–2430. PMLR.

Sid Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, Michael Pieler, USVSN Sai Prashanth, Shivanshu Purohit, Laria Reynolds, Jonathan Tow, Ben Wang, and Samuel Weinbach. 2022a. Gpt-neox-20b: An open-source autoregressive language model. *CoRR*, abs/2204.06745.

Sidney Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, and 1 others. 2022b. Gpt-neox-20b: An open-source autoregressive language model. In *Proceedings of BigScience Episode# 5–Workshop on Challenges & Perspectives in Creating Large Language Models*, pages 95–136.

Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, and Emma Brunskill. 2021. On the opportunities and risks of foundation models. *CoRR*, abs/2108.07258.

Łukasz Borchmann, Michał Pietruszka, Wojciech Jaśkowski, Dawid Jurkiewicz, Piotr Halama, Paweł Józiak, Łukasz Garncarek, Paweł Liskowski, Karolina Szyndler, Andrzej Gretkowski, and 1 others. 2025. Arctic-tilt. business document understanding at sub-billion scale. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 6: Industry Track)*, pages 264–283.

Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, and 1 others. 2022. Improving language models by retrieving from trillions of tokens. In *International conference on machine learning*, pages 2206–2240. PMLR.

Eugene Borisov and Nikolay Mikhaylovskiy. 2023. Team ntr@ automin 2023: Dolly llm improves minuting performance, semantic segmentation doesn't. In *Proceedings of the 16th International Natural Language Generation Conference: Generation Challenges*, pages 132–137.

Geoffrey C Bowker and Susan Leigh Star. 2000. *Sorting things out: Classification and its consequences*. MIT press.

Ivan Bulatovic, Sanja Škorić, and Vladimir Jovanović. 2016. Branding a business name. *Economics of agriculture*, 63(4):1323–1332.

Charlie Bullock, Suzanne Van Arsdale, Mackenzie Arnold, Cullen O'Keefe, and Christoph Winter. 2024. Legal considerations for defining "frontier model"'. *Institute for Law and AI*. Online; Available at: https://law-ai.org/frontier-model-definitions/.

Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, Xiaoyi Dong, Haodong Duan, Qi Fan, Zhaoye Fei, Yang Gao, Jiaye Ge, Chenya Gu, Yuzhe Gu, Tao Gui, and 81 others. 2024. Internlm2 technical report. *Preprint*, arXiv:2403.17297.

Anthony Cardillo. 2025. Number of Grok Users (Grok Statistics 2025). *Exploding Topics*. Online; Available at: https://explodingtopics.com/blog/grok-users.

Allan KK Chan and Yue-Yuan Huang. 2001. Chinese brand naming: a linguistic analysis of the brands of ten product categories. *Journal of product & brand Management*, 10(2):103–119.

ZhihaiTu Chat. 2025. ZhihaiTu Chat-ZhihaiTu Chat is an intelligent conversational assistant powered by a large language model under the Zhihu platform. — aibase.com. *Alibaba*. Online; Available at: https://www.aibase.com/tool/31346.

Rhitu Chatterjee. 2025. Their teenage sons died by suicide. now, they are sounding an alarm about AI chatbots. *NPR-American University Radio*. Online; Available at: https://www.npr.org/sections/shots-health-news/2025/09/19/nx-s1-5545749/ai-chatbots-safety-openai-meta-characterai-teens-suicide.

Gongwei Chen, Leyang Shen, Rui Shao, Xiang Deng, and Liqiang Nie. 2024. Lion: Empowering multimodal large language model with dual-level visual knowledge. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Jiuhai Chen, Zhiyang Xu, Xichen Pan, Yushi Hu, Can Qin, Tom Goldstein, Lifu Huang, Tianyi Zhou, Saining Xie, Silvio Savarese, Le Xue, Caiming Xiong, and Ran Xu. 2025. Blip3-o: A family of fully open unified multimodal models-architecture, training and dataset. *CoRR*, abs/2505.09568.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, and 39 others. 2021. Evaluating large language models trained on code. *Preprint*, arXiv:2107.03374.

Ye Chen, Wei Cai, Liangmin Wu, Xiaowei Li, Zhanxuan Xin, and Cong Fu. 2023a. Tigerbot: An open multilingual multitask llm. *arXiv preprint arXiv:2312.08688*.

Yirong Chen, Zhenyu Wang, Xiaofen Xing, Huimin Zheng, Zhipei Xu, Kai Fang, Junhong Wang, Sihang Li, Jieling Wu, Qi Liu, and 1 others. 2023b. Bianque: Balancing the questioning and suggestion ability of health llms with multi-turn health conversations polished by chatgpt. *CoRR*.

Zhihong Chen, Feng Jiang, Junying Chen, Tiannan Wang, Fei Yu, Guiming Chen, Hongbo Zhang, Juhao Liang, Chen Zhang, Zhiyi Zhang, and 1 others. 2023c. Phoenix: Democratizing chatgpt across languages. *arXiv preprint arXiv:2304.10453*.

Wei-Lin Chiang, Zhuohan Li, Ziqing Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, and 1 others. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. *See https://vicuna. lmsys. org (accessed 14 April 2023)*, 2(3):6.

Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.

Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, and 16 others. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53.

John Joon Young Chung and Max Kreminski. 2024. Patchview: Llm-powered worldbuilding with generative dust and magnet visualization. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*, pages 1–19.

Jack Clark. 2019. "Snuffleupagus, or Snuffy for short. We chose to name it GPT2 publicly as felt in poor taste to give muppet name while discussing reasons to be cautious with regard to increasingly powerful language models.". *Post on X (formerly Twitter)*. Online; Available at: https://x.com/jackclarkSF/status/1187824098916753408.

Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. ELECTRA: pre-training text encoders as discriminators rather than generators. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Team Cohere, Arash Ahmadian, Marwan Ahmed, Jay Alammar, Yazeed Alnumay, Sophia Althammer, Arkady Arkhangorodsky, Viraat Aryabumi, Dennis Aumiller, Rapha¨el Avalos, and 1 others. 2025. Command a: An enterprise-ready large language model. *arXiv preprint arXiv:2504.00698*.

Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, and Evan Rosen. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *Preprint*, arXiv:2507.06261.

Kate Crawford. 2021. *The atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

David Crystal. 2000. Muppet-Keep Your English Up to date. *BBC Article*. Online; Available at: https://www.bbc.co.uk/worldservice/learningenglish/radio/specials/1453_uptodate3/page14.shtml.

Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. Transformer-xl: Attentive language models beyond a fixed-length context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.

Yuhao Dan, Zhikai Lei, Yiyang Gu, Yong Li, Jianghao Yin, Jiaju Lin, Linhao Ye, Zhiyan Tie, Yougen Zhou, Yilei Wang, and 1 others. 2023. Educhat: A large-scale language model-based chatbot system for intelligent education. *arXiv preprint arXiv:2308.02773*.

Tri Dao and Albert Gu. 2024. Transformers are ssms: generalized models and efficient algorithms through structured state space duality. In *Proceedings of the 41st International Conference on Machine Learning*, pages 10041–10071.

DeepSeek. 2025. Deepseekai–Hugging Face. https://huggingface.co/deepseek-ai. [Accessed 14-05-2025].

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019a. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Association for Computational Linguistics*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019b. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Nolan Dey, Gurpreet Gosal, Hemant Khachane, William Marshall, Ribhu Pathria, Marvin Tom, Joel Hestness, and 1 others. 2023. Cerebras-gpt: Open compute-optimal language models trained on the cerebras wafer-scale cluster. *arXiv preprint arXiv:2304.03208*.

Disney Press Relations. 2004. The Walt Disney Company And The Jim Henson Company Sign Agreement For Disney To Buy The "Muppets" And "Bear In The Big Blue House". *The Walt Disney Company*. Online; Available at: https://thewaltdisneycompany.com/the-walt-disney-company-and-the-jim-henson-company-sign-agreement-for-disney-to-buy-the-muppets-and-bear-in-the-big-blue-house/.

Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. *Advances in neural information processing systems*, 32.

Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, and 3 others. 2023. Palm-e: An embodied multimodal language model. In *arXiv preprint arXiv:2303.03378*.

Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, and 1 others. 2022a. Glam: Efficient scaling of language models with mixture-of-experts. In *International Conference on Machine Learning*, pages 5547–5569. PMLR.

Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. 2022b. Glm: General language model pretraining with autoregressive blank infilling. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 320–335.

Maxim Enis and Mark Hopkins. 2024. From llm to nmt: Advancing low-resource machine translation with claude. *arXiv preprint arXiv:2404.13813*.

Muppet Fandom. 2025. Sesamstrasse. *Muppet Fandom Wiki*. Online; Available at: https://muppet.fandom.com/wiki/Sesamstrasse.

Jie Feng, Yuwei Du, Tianhui Liu, Siqi Guo, Yuming Lin, and Yong Li. 2024. Citygpt: Empowering urban spatial cognition of large language models. *arXiv preprint arXiv:2406.13948*.

Christopher Finch. 1993. *Jim Henson: The Works - The Art, the Magic, the Imagination*. Random House, New York, USA.

Andrew Mark Miller Fox News. 2026. Musk, xAI tout newest Grok update as only non-woke platform: Doesn't equivocate. Online; Available at: https://www.foxnews.com/politics/musk-xai-tout-newest-grok-update-as-only-non-woke-platform-citing-answers-to-key-questions.

James Frawley. 1979. The muppet movie. [Film]. Released on June 22, 1979, in the US.

Deep Ganguli, Amanda Askell, Nicholas Schiefer, Thomas I Liao, Kamilė Lukošiūtė, Anna Chen, Anna Goldie, Azalia Mirhoseini, Catherine Olsson, Danny Hernandez, and 1 others. 2023. The capacity for moral self-correction in large language models. *arXiv preprint arXiv:2302.07459*.

Burleigh B Gardner and Sidney J Levy. 1999. The product and the brand. *Sidney J. Levy und Dennis W. Rook (Hg.): Brands, consumers, symbols, & research. Sidney J. Levy on marketing. Thousand Oaks, Calif: Sage Publications*, pages 131–141.

Gregory J Gettas. 1990. The globalization of sesame street: A producer's perspective. *Educational Technology Research and Development*, 38(4):55–63.

Amelia Glaese, Nat McAleese, Maja Trębacz, John Aslanides, Vlad Firoiu, Timo Ewalds, Maribeth Rauh, Laura Weidinger, Martin Chadwick, Phoebe Thacker, and 1 others. 2022. Improving alignment of dialogue agents via targeted human judgements. *arXiv preprint arXiv:2209.14375*.

Maharshi Gor, Hal Daumé III, Tianyi Zhou, and Jordan Boyd-Graber. 2024. Do great minds think alike? investigating human-ai complementarity in question answering with caimira. In *Empirical Methods in Natural Language Processing*. Association for Computational Linguistics,.

Loren Graham. 2013. The power of names: In culture and in mathematics. *Proceedings of the American Philosophical Society*, 157(2):229–34.

Dirk Groeneveld, Iz Beltagy, Evan Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, and 1 others. 2024. Olmo: Accelerating the science of language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15789–15809.

Alexandre Grothendieck. 2021. *Récoltes et semailles (Reaping and Sowing),Part I: "Reflections and testimony on a past"*. Gallimard TEL.

Albert Gu and Tri Dao. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.

Fernando Guimarães. 1997. *Puppetry*. Pedia Press, Limburg, Germany.

Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Harkirat Singh Behl, Xin Wang, Sébastien Bubeck, Ronen Eldan, Adam Tauman Kalai, Yin Tat Lee, and Yuanzhi Li. 2023. Textbooks are all you need. *Preprint*, arXiv:2306.11644.

Nikhil Gupta and Jason Yip. 2024. Dbrx: Creating an llm from scratch using databricks. In *Databricks Data Intelligence Platform: Unlocking the GenAI Revolution*, pages 311–330. Springer.

Natalia El Hajj. 2024. Musk's Grok Chatbot Spreads US Election Results Misinformation. *insidetelecom.com*. Online; Available at: `https://insidetelecom.com/groks-ai-in-elections-misleads-american-voters/`.

Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. {DEBERTA}: {DECODING}-{enhanced} {bert} {with} {disentangled} {attention}. In *International Conference on Learning Representations*.

Brian Henson. 1992. The Muppet Christmas Carol. Movie by Walt Disney Pictures and Jim Henson Productions.

Jim Henson. 1983. Fraggle Rocck. Children's musical fantasy comedy puppet television series.

Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, and 1 others. 2022. Training compute-optimal large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, pages 30016–30030.

Huang Huang, Fei Yu, Jianqing Zhu, Xuening Sun, Hao Cheng, Dingjie Song, Zhihong Chen, Mosen Alharthi, Bang An, Juncai He, Ziche Liu, Junying Chen, Jianquan Li, Benyou Wang, Lian Zhang, Ruoyu Sun, Xiang Wan, Haizhou Li, and Jinchao Xu. 2024. Acegpt, localizing large language models in arabic. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), NAACL 2024, Mexico City, Mexico, June 16-21, 2024*, pages 8139–8163. Association for Computational Linguistics.

Kaixuan Huang, Xudong Guo, and Mengdi Wang. 2025. Specdec++: Boosting speculative decoding via adaptive candidate lengths. In *Second Conference on Language Modeling*.

Shaohan Huang, Li Dong, Wenhui Wang, Yaru Hao, Saksham Singhal, Shuming Ma, Tengchao Lv, Lei Cui, Owais Khan Mohammed, Barun Patra, Qiang Liu, Kriti Aggarwal, Zewen Chi, Johan Bjorck, Vishrav Chaudhary, Subhojit Som, Xia Song, and Furu Wei. 2023. Language is not all you need: Aligning perception with language models. *Preprint*, arXiv:2302.14045.

Yuqi Huo, Manli Zhang, Guangzhen Liu, Haoyu Lu, Yizhao Gao, Guoxing Yang, Jingyuan Wen, Heng Zhang, Baogui Xu, Weihao Zheng, Zongzheng Xi, Yueqian Yang, Anwen Hu, Jinming Zhao, Ruichen Li, Yida Zhao, Liang Zhang, Yuqing Song, Xin Hong, and 16 others. 2021. Wenlan: Bridging vision and language by large-scale multi-modal pre-training. *CoRR*, abs/2103.06561.

Srinivasan Iyer, Xi Victoria Lin, Ramakanth Pasunuru, Todor Mihaylov, Daniel Simig, Ping Yu, Kurt Shuster, Tianlu Wang, Qing Liu, Punit Singh Koura, Xian Li, Brian O'Horo, Gabriel Pereyra, Jeff Wang, Christopher Dewan, Asli Celikyilmaz, Luke Zettlemoyer, and Ves Stoyanov. 2023. Opt-iml: Scaling language model instruction meta learning through the lens of generalization. *Preprint*, arXiv:2212.12017.

Mohit Iyyer, Varun Manjunatha, Jordan Boyd-Graber, and Hal Daumé III. 2015. Deep unordered composition rivals syntactic methods for text classification. In *Proceedings of the Association for Computational Linguistics*.

Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2023. Atlas: Few-shot learning with retrieval augmented language models. *Journal of Machine Learning Research*, 24(251):1–43.

Zeke Jarvis. 2015. After the success of sam and friends, henson married his assistant from the project, jane nebel, in 1959. the couple had five children: Lisa, cheryl, brian, john, and heather. the family moved to new york city in 1963, where they set up the headquarters for muppets, inc. after jane quit working to raise the family. *Make'em Laugh! American Humorists of the 20th and 21st Centuries: American Humorists of the 20th and 21st Centuries*, page 138.

Mojan Javaheripi, Sébastien Bubeck, Marah Abdin, Jyoti Aneja, Sebastien Bubeck, Caio César Teodoro Mendes, Weizhu Chen, Allie Del Giorno, Ronen Eldan, Sivakanth Gopi, and 1 others. 2023. Phi-2: The surprising power of small language models. *Microsoft Research Blog*, 1(3):3.

Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego

de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, and 1 others. 2023a. Mistral 7b. *arXiv preprint arXiv:2310.06825*.

Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, Lélio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, and 7 others. 2024. Mixtral of experts. *Preprint*, arXiv:2401.04088.

Nan Jiang, Chengxiao Wang, Kevin Liu, Xiangzhe Xu, Lin Tan, Xiangyu Zhang, and Petr Babkin. 2025. Nova: Generative language models for assembly code with hierarchical attention and contrastive learning. In *The Thirteenth International Conference on Learning Representations*.

Wenxin Jiang, Chingwo Cheung, Mingyu Kim, Heesoo Kim, George K Thiruvathukal, and James C Davis. 2023b. Naming practices of pre-trained models in hugging face. *arXiv preprint arXiv:2310.01642*.

Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. The state and fate of linguistic diversity and inclusion in the NLP world. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 6282–6293. Association for Computational Linguistics.

Mari Kaestle, Will Morrison, John E. Barrett, Nancy Moran, and Rick Wetzel. 1980. *The Art of the Muppets*. Muppet Press/Bantam Books, New York, USA.

Yuki Kataoka, Tomohiro Takayama, Keisuke Yoshimura, Ryuhei So, Yasushi Tsujimoto, Yosuke Yamagishi, Shiro Takagi, Yuki Furukawa, Masatsugu Sakata, Đorđe Bašić, and 1 others. 2025. Automating the data extraction process for systematic reviews using gpt-4o and o3. *Research Synthesis Methods*, pages 1–21.

Kevin Lane Keller. 2013. Strategic brand managment: Building, measuring, and managing brand equity. *UK: Pearson Education Limited*.

John F. Kennedy. 2025. Address of Senator John F. Kennedy Accepting the Democratic Party Nomination for the Presidency of the United States, Los Angeles, CA, July 15, 1960. *John F. Kennedy Presidential Library*. Online; Available at: https://www.jfklibrary.org/archives/other-resources/john-f-kennedy-speeches/democratic-party-nomination-19600715.

DeAnna Kerley. 2013. The Names of 34 International Sesame Street Co-Productions. *Mental Floss*. Online; Available at: https://www.mentalfloss.com/article/53482/names-34-international-sesame-street-co-productions.

Omar Khattab and Matei Zaharia. 2020. Colbert: Efficient and effective passage search via contextualized late interaction over bert. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 39–48.

Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. 2020. Reformer: The efficient transformer. In *International Conference on Learning Representations*.

Philip Kottler and Kevin Lane Keller. 2009. Marketing management. *Jakarta: Erlangga*.

Zeljko Kraljevic, Anthony Shek, Daniel Bean, Rebecca Bendayan, James T. Teo, and Richard J. B. Dobson. 2021. Medgpt: Medical concept prediction from clinical narratives. *CoRR*, abs/2107.03134.

Zhenzhong Lan. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.

Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. Albert: A lite bert for self-supervised learning of language representations. In *International Conference on Learning Representations*.

J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.

Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, and 1 others. 2023. Bloom: A 176b-parameter open-access multilingual language model.

John D Lee and Katrina A See. 2004. Trust in automation: Designing for appropriate reliance. *Human factors*, 46(1):50–80.

Jan Leike, David Krueger, Tom Everitt, Miljan Martic, Vishal Maini, and Shane Legg. 2018. Scalable agent alignment via reward modeling: a research direction. *arXiv preprint arXiv:1811.07871*.

Jan Leike, John Schulman, and Jeffrey Wu. 2022. Our approach to alignment research. *OpenAI*. Online; Available at: https://openai.com/index/our-approach-to-alignment-research/.

Marco Lemos, Pedro JS Cardoso, and João MF Rodrigues. 2024. Harnessing ai and nlp tools for innovating brand name generation and evaluation: A comprehensive review. *Multimodal Technologies and Interaction*, 8(7):56.

Barak Lenz, Opher Lieber, Alan Arazi, Amir Bergman, Avshalom Manevich, Barak Peleg, Ben Aviram, Chen Almagor, Clara Fridman, Dan Padnos, and 1 others. 2025. Jamba: Hybrid transformer-mamba language models. In *The Thirteenth International Conference on Learning Representations*.

Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. 2021. {GS}hard: Scaling giant models with conditional computation and automatic sharding. In *International Conference on Learning Representations*.

Mike Lewis, Marjan Ghazvininejad, Gargi Ghosh, Armen Aghajanyan, Sida Wang, and Luke Zettlemoyer. 2020a. Pre-training via paraphrasing. In *Advances in Neural Information Processing Systems*, volume 33, pages 18470–18481. Curran Associates, Inc.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020b. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, and 1 others. 2020c. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:9459–9474.

Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. 2023. BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 19730–19742. PMLR.

Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.

Zhimin Li, Jianwei Zhang, Qin Lin, Jiangfeng Xiong, Yanxin Long, Xinchi Deng, Yingfang Zhang, Xingchao Liu, Minbin Huang, Zedong Xiao, Dayou Chen, Jiajun He, Jiahao Li, Wenyue Li, Chen Zhang, Rongwei Quan, Jianxiang Lu, Jiabin Huang, Xiaoyan Yuan, and 26 others. 2024. Hunyuan-dit: A powerful multi-resolution diffusion transformer with fine-grained chinese understanding. *Preprint*, arXiv:2405.08748.

Opher Lieber, Or Sharir, Barak Lenz, and Yoav Shoham. 2021. Jurassic-1: Technical details and evaluation. *White Paper. AI21 Labs*, 1(9):1–17.

Aixin Liu, Aoxue Mei, Bangcai Lin, Bing Xue, Bingxuan Wang, Bingzheng Xu, Bochao Wu, Bowei Zhang, Chaofan Lin, Chen Dong, and 1 others. 2025. Deepseek-v3. 2: Pushing the frontier of open large language models. *arXiv preprint arXiv:2512.02556*.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019a. Roberta: A robustly optimized bert pretraining approach. *Preprint*, arXiv:1907.11692.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Holistic Intelligencefor Global Good LLC. 2025. SeeChat x Ideas. Online; Available at: `https://seechat.ai/`.

Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32.

Kaixuan Lu, Ruiqian Zhang, Xiao Huang, and Yuxing Xie. 2024. Aquila: A hierarchically aligned visual-language model for enhanced remote sensing image comprehension. *CoRR*, abs/2411.06074.

Pan Lu, Baolin Peng, Hao Cheng, Michel Galley, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, and Jianfeng Gao. 2023. Chameleon: Plug-and-play compositional reasoning with large language models. *Advances in Neural Information Processing Systems*, 36:43447–43478.

Wei Lu, Rachel K Luu, and Markus J Buehler. 2025. Fine-tuning large language models for domain adaptation: Exploration of training strategies, scaling, model merging and synergistic capabilities. *npj Computational Materials*, 11(1):84.

Renqian Luo, Liai Sun, Yingce Xia, Tao Qin, Sheng Zhang, Hoifung Poon, and Tie-Yan Liu. 2022. Biogpt: generative pre-trained transformer for biomedical text generation and mining. *Briefings in bioinformatics*, 23(6):bbac409.

Zixuan Ma, Jiaao He, Jiezhong Qiu, Huanqi Cao, Yuanwei Wang, Zhenbo Sun, Liyan Zheng, Haojie Wang, Shizhi Tang, Tianyu Zheng, and 1 others. 2022. Bagualu: targeting brain scale pretrained models with over 37 million cores. In *Proceedings of the 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 192–204.

Helen McLure. 2000. The wild, wild web: The mythic american west and the electronic frontier. *Western Historical Quarterly*, 31(4):457–476.

Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love,

Pouya Tafti, and 1 others. 2024. Gemma: Open models based on gemini research and technology. *CoRR*.

Tomas Mikolov, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *International Conference on Learning Representations*.

Mayank Mishra, Matt Stallone, Gaoyuan Zhang, Yikang Shen, Aditya Prasad, Adriana Meza Soria, Michele Merler, Parameswaran Selvam, Saptha Surendran, Shivdeep Singh, and 1 others. 2024. Granite code models: A family of open foundation models for code intelligence. *CoRR*.

Thomas Mitchell. 2025. Trust and transparency in artificial intelligence: T. mitchell. *Philosophy & Technology*, 38(3):87.

Arindam Mitra, Luciano Del Corro, Shweti Mahajan, Andres Codas, Clarisse Simoes, Sahaj Agarwal, Xuxi Chen, Anastasia Razdaibiedina, Erik Jones, Kriti Aggarwal, and 1 others. 2023. Orca 2: Teaching small language models how to reason. *arXiv preprint arXiv:2311.11045*.

Niklas Muennighoff, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Jacob Morrison, Sewon Min, Weijia Shi, Pete Walsh, Oyvind Tafjord, Nathan Lambert, and 1 others. 2024. Olmoe: Open mixture-of-experts language models. *CoRR*.

Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and Ahmed Awadallah. 2023. Orca: Progressive learning from complex explanation traces of gpt-4. *arXiv preprint arXiv:2306.02707*.

Phoebe Mulcaire, Jungo Kasai, and Noah A Smith. 2019. Polyglot contextual representations improve crosslingual transfer. *arXiv preprint arXiv:1902.09697*.

muppetshow. 1976. Breaking the fourth wall in muppet films. *Muppet Fandom Wiki*. Online; Available at: https://muppet.fandom.com/wiki/Breaking_the_fourth_wall_in_Muppet_films.

Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, and 1 others. 2021. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*.

Marc Napolitano. 2009. Disneyfying dickens:" oliver & company" and" the muppet christmas carol" as dickensian musicals. *Studies in popular culture*, 32(1):79–102.

Sander Noels, Guillaume Bied, Maarten Buyl, Alexander Rogiers, Yousra Fettach, Jefrey Lijffijt, and Tijl De Bie. 2025. What large language models do not talk about: An empirical study of moderation and censorship practices. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 265–281. Springer.

OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, and Sam Altman. 2024. Gpt-4 technical report. *Preprint*, arXiv:2303.08774.

Riccardo Orlando, Luca Moroni, Pere-Lluís Huguet Cabot, Simone Conia, Edoardo Barba, Sergio Orlandini, Giuseppe Fiameni, and Roberto Navigli. 2024. Minerva llms: The first family of large language models trained from scratch on italian data. In *Proceedings of the 10th Italian Conference on Computational Linguistics (CLiC-it 2024)*, pages 707–719.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.

Arka Pal, Deep Karkhanis, Manley Roberts, Samuel Dooley, Arvind Sundararajan, and Siddartha Naidu. 2023. Giraffe: Adventures in expanding context lengths in llms. *arXiv preprint arXiv:2308.10882*.

Jupinder Parmar, Shrimai Prabhumoye, Joseph Jennings, Mostofa Patwary, Sandeep Subramanian, Dan Su, Chen Zhu, Deepak Narayanan, Aastha Jhunjhunwala, Ayush Dattagupta, and 1 others. 2024. Nemotron-4 15b technical report. *arXiv preprint arXiv:2402.16819*.

Roberta Pearson, Maire Messenger Davies, and Patrick Stewart. 2014. *Introduction: "It's a Television Show"*, 1 edition, pages 1–16. University of California Press.

Bo Peng, Eric Alcaide, Quentin Anthony, Alon Albalak, Samuel Arcadinho, Stella Biderman, Huanqi Cao, Xin Cheng, Michael Chung, Leon Derczynski, and 1 others. 2023. Rwkv: Reinventing rnns for the transformer era. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14048–14077.

Sarah Perkel. 2025. Sam altman says openai deserves to be mocked for its confusing ai names — and a 'fix' is coming. *msn.com-Business Insider*. Online; Available at: https://www.msn.com/en-us/money/companies/sam-altman-says-openai-deserves-to-be-mocked-for-its-confusing-ai-names-and-a-fix-is-coming/ar-AA1CYvgG.

Matthew E. Peters, Mark Neumann, Miran Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, Igor Livshits, and Mark Collins. 2018a. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.

Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018b. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for*

Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers), pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.

Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, and 1 others. 2018. Improving language understanding by generative pre-training.

Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.

Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, and 1 others. 2021. Scaling language models: Methods, analysis & insights from training gopher. *arXiv preprint arXiv:2112.11446*.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020a. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020b. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67.

Bal Ram and Pratima Verma. 2023. Artificial intelligence ai-based chatbot study of chatgpt, google ai bard and baidu ai. *World Journal of Advanced Engineering Technology and Sciences*, 8(01):258–261.

Neil Rathi, Dan Jurafsky, and Kaitlyn Zhou. 2025. Humans overrely on overconfident language models, across languages. *arXiv preprint arXiv:2507.06306*.

LG Research, Soyoung An, Kyunghoon Bae, Eunbi Choi, Stanley Jungkyu Choi, Yemuk Choi, Seokhee Hong, Yeonjung Hong, Junwon Hwang, Hyojin Jeon, and 1 others. 2024. Exaone 3.0 7.8 b instruction tuned language model. *arXiv preprint arXiv:2408.03541*.

Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, and 1 others. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*.

SAIC. 2012. Measures for the Registration of Enterprise Names by State Administration for Industry and Commerce. *State Administration for Industry and Commerce*. Online; Available at: https://www.eusmecentre.org.cn/sites/default/files/attach/Measures%20for%20Implementing%20the%20Administration%20of%20Enterprise%20Name%20Registration.pdf.

Salesforce. 2025. Einstein-platform. *Hugging Face*. Online; Available at: https://github.com/salesforce/einstein-platform.

Kishore Sampath, Ayaazuddin Mohammad, Resmi Ramachandranpillai, and 1 others. 2025. The multimodal paradox: How added and missing modalities shape bias and performance in multimodal ai. *arXiv preprint arXiv:2505.03020*.

Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.

Victor Sanh, Albert Webson, Colin Raffel, Stephen H Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, and 1 others. 2022. Multitask prompted training enables zero-shot task generalization. In *ICLR 2022-Tenth International Conference on Learning Representations*.

Alexander Scarlatos and Andrew S. Lan. 2023. Tree-based representation and generation of natural and mathematical language. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 3714–3730. Association for Computational Linguistics.

Jonathan Schroeder, Janet Borgerson, and Zhiyan Wu. 2015. A brand culture approach to chinese cultural heritage brands. *Journal of Brand Management*, 22(3):261–279.

Neha Sengupta, Sunil Kumar Sahu, Bokang Jia, Satheesh Katipomu, Haonan Li, Fajri Koto, William Marshall, Gurpreet Gosal, Cynthia Liu, Zhiming Chen, and 1 others. 2023. Jais and jais-chat: Arabic-centric foundation and instruction-tuned open generative large language models. *arXiv preprint arXiv:2308.16149*.

Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askell, Samuel R Bowman, Esin DURMUS, Zac Hatfield-Dodds, Scott R Johnston, Shauna M Kravec, and 1 others. 2023. Towards understanding sycophancy in language models. In *The Twelfth International Conference on Learning Representations*.

Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2019. Megatron-lm: Training multi-billion parameter language models using model parallelism. *ArXiv*, abs/1909.08053.

Sebastian Sigler. 2014. Die allwissende müllhalde. *Die Börse am Sonntag*. Online; Available at: https://www.boerse-am-sonntag.de/maerkte/die-allwissende-muellhalde.

Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhumoye, George

Zerveas, Vijay Korthikanti, Elton Zhang, Rewon Child, Reza Yazdani Aminabadi, Julie Bernauer, Xia Song, Mohammad Shoeybi, Yuxiong He, Michael Houston, Saurabh Tiwary, and Bryan Catanzaro. 2022. Using deepspeed and megatron to train megatron-turing nlg 530b, a large-scale generative language model. *Preprint*, arXiv:2201.11990.

S Sorensen. Claude 4.0, with manus ai (2025).". *Consciousness Assessment in Large Language Models: A Comparative Analysis of Response Patterns to Recursive Self-Examination and Temporal Discontinuity." Human-AI Collaborative Research Series*, 1(1).

Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2020. Energy and policy considerations for modern deep learning research. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 13693–13696.

Xingwu Sun, Yanfeng Chen, Yiqing Huang, Ruobing Xie, Jiaqi Zhu, Kai Zhang, Shuaipeng Li, Zhen Yang, Jonny Han, Xiaobo Shu, Jiahao Bu, Zhongzhi Chen, Xuemeng Huang, Fengzong Lian, Saiyong Yang, Jianfeng Yan, Yuyuan Zeng, Xiaoqin Ren, Chao Yu, and 89 others. 2024. Hunyuan-large: An open-source moe model with 52 billion activated parameters by tencent. *Preprint*, arXiv:2411.02265.

Zhiqing Sun, Yikang Shen, Qinhong Zhou, Hongxin Zhang, Zhenfang Chen, David Cox, Yiming Yang, and Chuang Gan. 2023. Principle-driven self-alignment of language models from scratch with minimal human supervision. *Advances in Neural Information Processing Systems*, 36:2511–2565.

Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. Alpaca: A strong, replicable instruction-following model. *Stanford Center for Research on Foundation Models. https://crfm. stanford. edu/2023/03/13/alpaca. html*, 3(6):7.

Yi Tay, Mostafa Dehghani, Vinh Q. Tran, Xavier Garcia, Jason Wei, Xuezhi Wang, Hyung Won Chung, Dara Bahri, Tal Schuster, Steven Zheng, Denny Zhou, Neil Houlsby, and Donald Metzler. 2023. UL2: Unifying language learning paradigms. In *The Eleventh International Conference on Learning Representations*.

Ross Taylor, Marcin Kardas, Guillem Cucurull, Thomas Scialom, Anthony Hartshorn, Elvis Saravia, Andrew Poulton, Viktor Kerkez, and Robert Stojnic. 2022. Galactica: A large language model for science. *arXiv preprint arXiv:2211.09085*.

Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, and 1 others. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.

HunyuanWorld Team, Zhenwei Wang, Yuhao Liu, Junta Wu, Zixiao Gu, Haoyuan Wang, Xuhui Zuo, Tianyu Huang, Wenhuan Li, Sheng Zhang, Yihang Lian, Yulin Tsai, Lifu Wang, Sicong Liu, Puhua Jiang, Xianghui Yang, Dongyuan Guo, Yixuan Tang, Xinyue Mao, and 36 others. 2025. Hunyuanworld 1.0: Generating immersive, explorable, and interactive 3d worlds from words or pixels. *CoRR*, abs/2507.21809.

Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, and 1 others. 2022. Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*.

Yuanhe Tian, Ruyi Gan, Yan Song, Jiaxing Zhang, and Yongdong Zhang. 2024. Chimed-gpt: A chinese medical large language model with full training regime and better alignment to human preferences. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7156–7173.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. Llama: Open and efficient foundation language models. *Preprint*, arXiv:2302.13971.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023b. Llama 2: Open foundation and fine-tuned chat models. *Preprint*, arXiv:2307.09288.

Lewis Tunstall, Edward Emanuel Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro Von Werra, Clémentine Fourrier, Nathan Habib, Nathan Sarrazin, Omar Sanseviero, Alexander M Rush, and Thomas Wolf. 2024. Zephyr: Direct distillation of LM alignment. In *First Conference on Language Modeling*.

Frederick Jackson Turner. 1893. The significance of the frontier in american history. In *The Structure of Political Geography*, 2017 edition, pages 132–139. Routledge.

A Vaswani. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.

James Vincent. 2019. Why are so many AI systems named after Muppets? Online; Available at: `https://www.theverge.com/2019/12/11/20993407/ai-language-models-muppets-sesame-street-muppetware-elmo-bert-ernie`.

Patrick Vonplaten. 2025. MistralAI. *Hugging Face*. Online; Available at: `https://huggingface.co/mistralai`.

Hannah Waight, Eddie Yang, Yin Yuan, Solomon Messing, Margaret E Roberts, Brandon M Stewart, and

Joshua A Tucker. 2026. State media control influences large language models. *Nature*, pages 1–9.

Ben Wang and Aran Komatsuzaki. 2021. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. https://github.com/kingoflolz/mesh-transformer-jax.

Guan Wang, Sijie Cheng, Xianyuan Zhan, Xiangang Li, Sen Song, and Yang Liu. 2024a. Openchat: Advancing open-source language models with mixed-quality data. In *ICLR*.

Shansong Wang, Mingzhe Hu, Qiang Li, Mojtaba Safari, and Xiaofeng Yang. 2025. Capabilities of gpt-5 on multimodal medical reasoning. *arXiv preprint arXiv:2508.08224*.

Zihan Wang, Liuxz2@chinatelecom.cn Liuxz2@chinatelecom.cn, Liusx14@chinatelecom.cn Liusx14@chinatelecom.cn, Yitong Yao, Huangyy121@chinatelecom.cn Huangyy121@chinatelecom.cn, Li Mengxiang, Zhongjiang He, Liyx25@chinatelecom.cn Liyx25@chinatelecom.cn, Pulw@chinatelecom.cn Pulw@chinatelecom.cn, Xuhn@chinatelecom.cn Xuhn@chinatelecom.cn, Chao Wang, and Shuangyong Song. 2024b. TeleChat: An open-source billingual large language model. In *Proceedings of the 10th SIGHAN Workshop on Chinese Language Processing (SIGHAN-10)*, pages 10–20, Bangkok, Thailand. Association for Computational Linguistics.

Zihan Wang, Yitong Yao, Li Mengxiang, Zhongjiang He, Chao Wang, Shuangyong Song, and 1 others. 2024c. Telechat: An open-source billingual large language model. In *Proceedings of the 10th SIGHAN Workshop on Chinese Language Processing (SIGHAN-10)*, pages 10–20.

Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. 2022. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*.

Jerry Wei, Da Huang, Yifeng Lu, Denny Zhou, and Quoc V Le. 2023. Simple synthetic data reduces sycophancy in large language models. *arXiv preprint arXiv:2308.03958*.

Laura Weidinger, Jonathan Uesato, Jakub Bielecki, Glenn van den Driessche, Mike Chrzanowski, Dmitriy Krasheninnikov, Martin Chadwick, Rohen Shah Gur, Amanda Glaese, Ruben Tréger, and 1 others. 2021. Ethical and social risks of large language models. *arXiv preprint arXiv:2112.04359*.

Jinting Wu. 2016. Educational discipline, ritual governing, and chinese exemplary society: Why china's curriculum reform remains a difficult task. *Policy Futures in Education*, 14(6):721–740.

Shaohua Wu, Xudong Zhao, Shenling Wang, Jiangang Luo, Lingjun Li, Xi Chen, Bing Zhao, Wei Wang, Tong Yu, Rongguo Zhang, Jiahua Zhang, and Chao Wang. 2023a. Yuan 2.0: A large language model with localized filtering-based attention. *Preprint*, arXiv:2311.15786.

Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhanjan Kambadur, David Rosenberg, and Gideon Mann. 2023b. Bloomberggpt: A large language model for finance. *arXiv preprint arXiv:2303.17564*.

Yongliang Wu, Yizhou Zhou, Zhou Ziheng, Yingzhe Peng, Xinyu Ye, Xinting Hu, Wenbo Zhu, Lu Qi, Ming-Hsuan Yang, and Xu Yang. 2025. On the generalization of sft: A reinforcement learning perspective with reward rectification. *arXiv preprint arXiv:2508.05629*.

xAI. 2025. Grok. *Hugging Face*. Online; Available at: https://github.com/xai-org/grok-1.

Xi Jinping. 2025. Speech at the forum on literature and art. *Friends of Socialist China*. Online; Available at: https://socialistchina.org/2025/03/04/xi-jinping-speech-at-the-forum-on-literature-and-art/.

Canwen Xu, Daya Guo, Nan Duan, and Julian McAuley. 2023. Baize: An open-source chat model with parameter-efficient tuning on self-chat data. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6268–6278.

Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. mT5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

Aiyuan Yang, Bin Xiao, Bingning Wang, Borong Zhang, Ce Bian, Chao Yin, Chenxu Lv, Da Pan, Dian Wang, Dong Yan, and 1 others. 2023. Baichuan 2: Open large-scale language models. *arXiv preprint arXiv:2309.10305*.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 40 others. 2025a. Qwen3 technical report. *CoRR*, abs/2505.09388.

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, and 1 others. 2024. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.

Jianwei Yang, Reuben Tan, Qianhui Wu, Ruijie Zheng, Baolin Peng, Yongyuan Liang, Yu Gu, Mu Cai, Seonghyeon Ye, Joel Jang, and 1 others. 2025b. Magma: A foundation model for multimodal ai agents. In *Proceedings of the Computer Vision*

and Pattern Recognition Conference, pages 14203–14214.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32.

Xinman Yu. 2023. A study on the selection and use of classical literature in partially compiled high school language textbooks. *Advances in Higher Education*, 7(3):129–133.

Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2020a. Big bird: Transformers for longer sequences. In *Advances in Neural Information Processing Systems*, volume 33, pages 17283–17297. Curran Associates, Inc.

Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and 1 others. 2020b. Big bird: Transformers for longer sequences. *Advances in neural information processing systems*, 33:17283–17297.

Fabio Massimo Zanzotto, Andrea Santilli, Leonardo Ranaldi, Dario Onorati, Pierfrancesco Tommasino, and Francesca Fallucchi. 2020. KERMIT: Complementing transformer architectures with encoders of explicit syntactic interpretations. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 256–267, Online. Association for Computational Linguistics.

Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019a. Defending against neural fake news. *Advances in neural information processing systems*, 32.

Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019b. Defending against neural fake news. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

Aohan Zeng, Xin Lv, Qinkai Zheng, Zhenyu Hou, Bin Chen, Chengxing Xie, Cunxiang Wang, Da Yin, Hao Zeng, Jiajie Zhang, and Kedong Wang and. 2025. GLM-4.5: agentic, reasoning, and coding (ARC) foundation models. *CoRR*, abs/2508.06471.

Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Diego Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, Hao Yu, Hongning Wang, Jiadai Sun, Jiajie Zhang, Jiale Cheng, Jiayi Gui, Jie Tang, Jing Zhang, Juanzi Li, Lei Zhao, and 36 others. 2024. Chatglm: A family of large language models from GLM-130B to GLM-4 all tools. *CoRR*, abs/2406.12793.

Wei Zeng, Xiaozhe Ren, Teng Su, Hui Wang, Yi Liao, Zhiwei Wang, Xin Jiang, ZhenZhang Yang, Kaisheng Wang, Xiaoda Zhang, and 1 others. 2021. Pangu-α: Large-scale autoregressive pretrained chinese language models with auto-parallel computation. *arXiv preprint arXiv:2104.12369*.

Hongbo Zhang, Junying Chen, Feng Jiang, Fei Yu, Zhihong Chen, Guiming Chen, Jianquan Li, Xiangbo Wu, Zhiyi Zhang, Qingying Xiao, Xiang Wan, Benyou Wang, and Haizhou Li. 2023. Huatuogpt, towards taming language model to be a doctor. In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 10859–10885. Association for Computational Linguistics.

Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020a. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International conference on machine learning*, pages 11328–11339. PMLR.

Kaiqi Zhang, Jing Zhao, and Rui Chen. 2024. Koala: Enhancing speculative decoding for llm via multi-layer draft heads with adversarial learning. *arXiv preprint arXiv:2408.08146*.

Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, and 1 others. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.

Xuanyu Zhang and Qing Yang. 2023. Xuanyuan 2.0: A large chinese financial chat model with hundreds of billions parameters. In *Proceedings of the 32nd ACM international conference on information and knowledge management*, pages 4435–4439.

Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020b. Dialogpt: Large-scale generative pre-training for conversational response generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*. Association for Computational Linguistics.

Zhengyan Zhang, Yuxian Gu, Xu Han, Shengqi Chen, Chaojun Xiao, Zhenbo Sun, Yuan Yao, Fanchao Qi, Jian Guan, Pei Ke, and 1 others. 2021a. Cpm-2: Large-scale cost-effective pre-trained language models. *arXiv preprint arXiv:2106.10715*.

Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019a. Ernie: Enhanced language representation with informative entities. *arXiv preprint arXiv:1905.07129*.

Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019b. ERNIE: Enhanced language representation with informative entities. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1441–1451, Florence, Italy. Association for Computational Linguistics.

Zhuosheng Zhang, Hanqing Zhang, Keming Chen, Yuhang Guo, Jingyun Hua, Yulong Wang, and Ming Zhou. 2021b. Mengzi: Towards lightweight yet ingenious pre-trained models for chinese. *Preprint*, arXiv:2110.06696.

Tianyang Zhong, Zhengliang Liu, Yi Pan, Yutong Zhang, Yifan Zhou, Shizhe Liang, Zihao Wu, Yanjun Lyu, Peng Shu, Xiaowei Yu, and 1 others. 2024. Evaluation of openai o1: Opportunities and challenges of agi. *CoRR*.

Hao Zhou, Pei Ke, Zheng Zhang, Yuxian Gu, Yinhe Zheng, Chujie Zheng, Yida Wang, Chen Henry Wu, Hao Sun, Xiaocong Yang, Bosi Wen, Xiaoyan Zhu, Minlie Huang, and Jie Tang. 2021. Eva: An open-domain chinese dialogue system with large-scale generative pre-training. *Preprint*, arXiv:2108.01547.

Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2019. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*.

Jack Zipes. 2015. *The Oxford companion to fairy tales*. Oxford University Press (UK).

A Moopet Models

Indeed, probably the best reason against referring MMs as Muppet Models is that the term "Muppet" is owned by the Disney Corporation. Given Disney's ownership of the Muppets, it is uncertain how receptive the company would be to such a designation. Not being an IP lawyer, professional advice cannot be provided, but considering the extensive use of characters like Elmo and BERT, it may well be that the opportunity for uncontested use of the name "Muppet Models" has passed.

However, the Muppet Movie (Frawley, 1979) may present another alternative. Fozzie joins a group called the Moopets, which although similar in appearance to the Muppets, seems to be sufficiently legally distinct from the actual Muppets to evade legal scrutiny in the Muppet universe. While this academic paper does not dare to give legal advice, calling models "Moopet models" might be a legal grey zone that would convey the essence of these ideas while not causing confusion for the public.

B General Naming Criteria

Model names are not neutral designations; they function as mechanisms of branding, cultural alignment, technical signaling, and market positioning. Across the dataset, we identify several recurring naming strategies.

- **Muppet-inspired acronyms.** Early NLP models frequently drew inspiration from *Sesame Street* and related Muppet characters, including ELMo (Peters et al., 2018a), BERT (Devlin et al., 2019b), ERNIE, KERMIT (Zanzotto et al., 2020), BigBird, and Rosita. Although playful, this practice established a recognizable naming tradition that reflected the academic and community-oriented culture of late-2010s NLP research.

  The trend began with ELMo ("Embeddings from Language Models"), whose acronym intentionally evokes the *Sesame Street* character Elmo. Its influence expanded with BERT ("Bidirectional Encoder Representations from Transformers"), introduced by Google in 2018. Although Google never formally confirmed the reference, the naming pattern strongly suggested continuity with ELMo and the broader Muppet-inspired convention (Vincent, 2019). The success of BERT encouraged subsequent researchers to adopt similar naming schemes, leading to models such as ERNIE and KERMIT. OpenAI researchers also reportedly considered the name "Snuffleupagus" during GPT's early development (Vincent, 2019). Over time, these examples established a coherent cultural naming pattern within NLP research.

  Muppet-inspired names provide several advantages. They are memorable, approachable, and culturally recognizable, helping technical systems appear less intimidating. Such names also facilitate dissemination across social media, academic discussions, and popular reporting by increasing recall and emotional resonance.

- **Capability-driven names.** Many models employ names that foreground architectural properties, capabilities, or intended functionality rather than cultural references. Examples include T5 ("Text-to-Text Transfer Transformer") (Raffel et al., 2020a), XLNet (Yang et al., 2019), UniLM, FLAN, OPT, UL2, LaMDA, Codex, OLMo, and DBRX.

  This naming strategy became especially prominent after GPT ("Generative Pre-trained Transformer"), which explicitly emphasized the model architecture rather than a character reference. Subsequent models extended this

approach through descriptive naming conventions that signal scale, efficiency, multimodality, dialogue capability, or instruction tuning. For example, XLNet emphasizes "eXtreme Large Transformer" capabilities, while T5 directly communicates its unified text-to-text learning framework.

Capability-driven names prioritize clarity and technical communication. They allow researchers and practitioners to infer functionality rapidly while emphasizing engineering objectives and architectural innovation.

- **Popular culture references.** Some models borrow from broader cultural or fictional repertoires, prioritizing memorability and branding over direct technical description. Examples include BART (Lewis et al., 2020b), Megatron-LM, Grok (xAI, 2025), Nemotron, and GPT-Neo.

  These names often evoke recognizable fictional characters or science-fiction traditions. BART references *The Simpsons*, while Megatron and Nemotron evoke the *Transformers* franchise. Grok derives from Robert Heinlein's science-fiction vocabulary, where "grok" signifies deep intuitive understanding. Such naming practices create distinctive and memorable identities while associating models with intelligence, scale, or futuristic imagination.

- **Derivations from existing models.** Many models communicate lineage by retaining an established base name and appending modifiers that indicate extensions, improvements, or variants. Examples include RoBERTa (Liu et al., 2019a), ALBERT (Lan, 2019), DistilBERT (Sanh et al., 2019), mT5, CPM-2, LLaMA-2, Code Llama, ChatGLM2, Mixtral 2, Claude 4.1, and Qwen 3.

  This strategy allows developers to signal both continuity and innovation simultaneously. For example, RoBERTa explicitly positions itself as an optimized extension of BERT, while ALBERT ("A Lite BERT") emphasizes parameter efficiency. DistilBERT foregrounds compression through knowledge distillation, and Code Llama indicates a code-specialized variant of LLaMA.

  Derivational naming conventions are especially common in rapidly evolving model families because they preserve brand recognition while clearly indicating architectural evolution.

- **Personal, historical, or scientific names.** Some models adopt human names or references to historical and scientific figures in order to evoke familiarity, authority, intelligence, or prestige. Examples include Claude, Einstein, and CAIMIRA.

  These names anthropomorphize models and make them appear more approachable while simultaneously invoking intellectual legitimacy. Claude is widely understood as a reference to Claude Shannon, while Einstein directly invokes scientific genius and innovation. Such naming practices blur the boundary between technical systems and human expertise, positioning models as knowledgeable collaborators rather than purely computational artifacts.

- **Mythological and historical references.** A common naming strategy draws from mythology, religion, or historical figures, often to signal wisdom, creativity, scale, or power. Examples include ELECTRA, Pegasus, Atlas, Phoenix, Pythia, Athene, Minerva, Mengzi, Pangu, Xuanyuan, Hou Yi, Jiutian, Tianshu, and Baize.

  This practice is especially prominent in Chinese models. For example, Pangu references the primordial creator figure in Chinese mythology, while Xuanyuan invokes the Yellow Emperor. Mengzi derives from the philosopher Mencius, associating the model with scholarship and moral reasoning. Western examples similarly invoke classical mythology, such as Pegasus, Atlas, Athene, and Minerva.

  Mythological and historical naming strategies associate AI systems with enduring cultural narratives, conveying authority, ambition, creativity, or wisdom through symbolic reference.

- **GPT + domain-specific names.** Another major naming category combines the GPT label with domain-specific modifiers, emphasizing specialization while preserving association with the GPT architecture. Examples include BioGPT, WebGPT, DialoGPT, Instruct-

GPT, BloombergGPT, HuatuoGPT, MathGPT, CityGPT, and MiniGPT-4.

These names communicate both architectural lineage and application domain. BioGPT emphasizes biomedical applications, BloombergGPT signals financial specialization, and HuatuoGPT associates the model with Chinese medical contexts. The naming strategy is particularly effective because it simultaneously leverages GPT's widespread recognition while differentiating models by task or industry.

- **Geographical names.** Some models adopt names associated with locations, landmarks, or geographical imagery. Examples include Xihu and Arctic.

  Geographical references are especially common in Chinese naming traditions, where locations frequently carry cultural, historical, or symbolic significance. Xihu ("West Lake") invokes one of China's most culturally significant landmarks, while Arctic evokes scale, extremity, and environmental imagery. Such names often function symbolically rather than literally, conveying prestige, breadth, or regional identity.

- **Company branding names.** Some models foreground institutional or corporate identity directly within the model name. Examples include DeepSeek, StableLM, SEEChat, TeleChat, AiLMe, and o1.

  These names emphasize organizational branding and strategic positioning rather than mythology or technical description. StableLM directly associates the model with Stability AI, while DeepSeek foregrounds company identity through a concise brand-oriented construction. Such naming strategies are especially common among commercial systems competing for recognition in increasingly crowded markets.

- **Flora and fauna names.** Animal- and plant-inspired names are another recurring strategy. Examples include LLaMA (Touvron et al., 2023a), Chinchilla, Sparrow, Falcon, Orca, Alpaca, Vicuna, Koala, Mamba, Chameleon, TigerBot, and Aquila.

  These names often combine memorability with metaphorical meaning. LLaMA simultaneously references the animal and abbreviates "Large Language Model Meta AI," while Falcon and Sparrow evoke speed, agility, or intelligence. Animal names can make models appear less abstract and more distinctive, while also providing flexible branding opportunities.

- **Traditional concepts.** Some models derive their names from longstanding philosophical, ethical, or cultural concepts rather than specific individuals or places. Examples include Wu Dao, Hunyuan, Tongyi Qianwen, WenLan, and Baichuan.

  These names often emphasize harmony, universality, knowledge, or holistic understanding. Wu Dao invokes the Daoist concept of "the Way," while Hunyuan suggests primordial unity and totality. Such naming strategies emphasize continuity with intellectual traditions and align AI development with broader cultural and philosophical frameworks.

- **Scientific and conceptual names.** Another class of models adopts names inspired by scientific phenomena, abstract concepts, or large-scale metaphors. Examples include Jurassic-1, Galactica, Kosmos-1, Mistral, Nova, Granite, Zephyr, Sonus, and Luminous.

  These names frequently evoke magnitude, exploration, dynamism, or sophistication. Galactica and Kosmos suggest scale and universality, while Zephyr and Mistral evoke speed and movement through meteorological imagery. Such naming conventions position models as technologically advanced, expansive, and conceptually ambitious.

C A Census of Language Model Names

Table 1 provides a curated inventory of large language models, categorizing them by naming origin, release year, regional affiliation, and associated publication or documentation. This census captures the diversity and evolution of naming conventions across both Western and Chinese AI ecosystems. By organizing models into categories such as Muppet-inspired names, capability-driven names, mythological references, flora and fauna names, and company branding names, the table serves both as a dataset and as an analytical foundation for understanding how naming practices reflect broader

institutional, cultural, technical, and commercial priorities in language model development.

Model Name The official public name assigned to the model at release. When applicable, both English and localized names are included. Examples include BERT, GPT, ERNIE, Tongyi Qianwen, and DeepSeek. Model names often encode technical lineage, branding strategy, cultural symbolism, or intended application domains.

Type This field records the naming category assigned to each model. Categories were manually curated based on the apparent linguistic, cultural, or strategic motivation behind the model name. The categories include:

- Muppet Character-Inspired Names
- Capability-Driven Names
- Popular Culture References
- Derivations from Existing Models
- Mythological & Historical References
- Personal, Historical, or Scientific Names
- GPT + Domain-Specific Names
- Flora and Fauna Names
- Traditional Concepts
- Scientific / Conceptual Names
- Geographical Names
- Company Branding Names
- English Names

These categories reflect the rationale, symbolism, or branding logic underlying each model name rather than the technical architecture itself.

Release Year The year in which the model was publicly introduced through a research paper, technical report, preprint, open-source release, product launch, or major public announcement.

Source The table distinguishes between models originating primarily from Chinese and Western research or industrial ecosystems. This distinction allows comparative analysis of naming conventions across cultural and institutional contexts.

Title This field contains either the official paper title, technical report title, or a short descriptive phrase associated with the model. These descriptions provide contextual information about the model's intended purpose, architectural contribution, or application domain.

Model name	Type	Year	Source	Title
BERT	Muppet Character-Inspired Names	2018	Western	Bidirectional Encoder Representations from Transformers
ELMO	Muppet Character-Inspired Names	2018	Western	Embeddings from Language Models
RoBERTA	Derivation from original	2018	Western	Robustly optimized BERT Pre-training Approach
DistilBERT	Derivation from original	2019	Western	a distilled version of BERT
ELECTRA	Mythological & Historical References	2019	Western	Efficient Pre-training via replaced token detection
ERNIE	Muppet Character-Inspired Names	2019	China	Enhanced Language Representation with Informative Entities
GPT	GPT + Domain-Specific Names	2019	Western	Generative pre-trained transformer
Grover	Muppet Character-Inspired Names	2019	Western	Generating aRticles by Only Viewing mEtadata Records
KERMIT	Muppet Character-Inspired Names	2019	Western	Kontextuell Encoder Representations Made by Insertion Transformation
Megatron-LM	Popular culture references	2019	Western	Megatron-LM
Rosita	Muppet Character-Inspired Names	2019	Western	Polyglot Contextual Representations Improve Crosslingual Transfer
T5	Capability-driven names	2019	Western	Text-to-Text Transfer Transformer
Trransformer-XL	Capability-driven names	2019	Western	Transformer-xl: Attentive language models beyond a fixed-length context
UniLM	Capability-driven names	2019	Western	Unified Language Model
XLNet	Capability-driven names	2019	Western	"eXtreme Large Transformer" combined with "Permutation Language Modeling"
ColBERT	Derivation from original	2020	Western	Contextualized Late Interaction over BERT
ALBERT	Derivation from original	2020	Western	Albert: A Lite Bert For Self-Supervised Learning Of Language Representations
BART	Popular culture references	2020	Western	Bidirectional and Autoregressive Transformer
Big BIRD	Muppet Character-Inspired Names	2020	Western	Big Bird: Transformers for Longer Sequences
DeBERTa	Derivation from original	2020	Western	Decoding-enhanced BERT with disentangled attention
DialoGPT	GPT + Domain-Specific Names	2020	Western	DIALOGPT: Large-Scale Generative Pre-training for Conversational Response Generation

GPT-Neo	Popular culture references	2020	Western	Generative Pre-trained Transformer Neo
MARGE	Muppet Character-Inspired Names	2020	Western	Multilingual Autoencoder that Retrieves and Generates
mT5	Derivation from original	2020	Western	multilingual variant of T5
Pegasus	Mythological & Historical References	2020	Western	Pegasus: Pre-training with extracted gap-sentences for abstractive summarization
Reformer	Capability-driven names	2020	Western	Reformer: The Efficient Transformer
Codex	Capability-driven names	2021	Western	Evaluating Large Language Models Trained on Code
CPM	Capability-driven names	2021	China	Chinese Pre-trained Model
CPM-2	Derivation from original	2021	China	Chinese Pre-trained Model 2
EVA	Capability-driven names	2021	China	Enhanced Viirtual Assistant
FLAN	Capability-driven names	2021	Western	Finetuned Language Net
GLaM	Capability-driven names	2021	Western	Generalist Language Model
GPT-J	GPT + Domain-Specific Names	2021	Western	GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model
GPT-NeoX-20B	Derivation from original	2021	Western	Gpt-neox-20b: An open-source autoregressive language model
GShard	Capability-driven names	2021	Western	Gshard: Scaling giant models with conditional computation and automatic sharding
Jurassic-1	Scientific / Conceptual Names	2021	Western	Jurassic-1: Large autoregressive language model by AI21 Labs
Mengzi	Mythological & Historical References	2021	China	Mengzi: Towards Lightweight yet Ingenious Pre-trained Models for Chinese
MT-NLG	Capability-driven names	2021	Western	Megatron-Turing NLG
Pangu	Mythological & Historical References	2021	China	Pangu Large-scale autoregressive pretrained Chinese language models with auto-parallel computation
WebGPT	GPT + Domain-Specific Names	2021	Western	WebGPT
WenLan	Traditional Concepts	2021	China	WenLan: Bridging Vision and Language by Large-Scale Multi-Modal Pre-Training
Bagualu	Mythological & Historical References	2022	China	BaGuaLu: targeting brain scale pretrained models with over 37 million cores
BioGPT	GPT + Domain-Specific Names	2022	Western	BioGPT: generative pre-trained transformer for biomedical text generation and mining

Chinchilla	Flora and Fauna Names	2022	Western	Chinchilla
Flan-PaLM	Capability-driven names	2022	Western	PaLM: Scaling Language Modeling with Pathways
Flan-T5	Capability-driven names	2022	Western	See FLAN and T5
Galactica	Scientific / Conceptual Names	2022	Western	Galactica
GLM-130B	Capability-driven names	2022	Western	General Language Model (GLM)
Gopher	Muppet Character-Inspired Names	2022	Western	Gopher
InstructGPT	GPT + Domain-Specific Names	2022	Western	InstructGPT
LaMDA	Capability-driven names	2022	Western	Language Model for Dialogue Applications
OPT	Capability-driven names	2022	Western	Open Pre-trained Transformer Language Models
OPT-IML	Capability-driven names	2022	Western	Open Pre-trained Transformer Instruction Meta Learning
Retro	Capability-driven names	2022	Western	Retrieval-Enhanced Transformer
Sparrow	Flora and Fauna Names	2022	Western	Sparrow
T0	Capability-driven names	2022	Western	T0
UL2	Capability-driven names	2022	Western	Unifying Language Learning Paradigms
AiLMe	Company branding names	2023	China	AiLMe
Alpaca	Flora and Fauna Names	2023	Western	Alpaca: A strong, replicable instruction-following model
Atlas	Mythological & Historical References	2023	Western	Atlas
AceGPT	GPT + Domain-Specific Names	2023	China	AceGPT: Localizing Large Language Models in Arabic
Aquila	Flora and Fauna Names	2023	China	Aquila series by BAAI
Baichuan	Traditional Concepts	2023	China	Baichuan: Open large-scale language models
Baichuan 2	Derivation from original	2023	China	Baichuan: Open large-scale language models
Baize	Mythological & Historical References	2023	Western	Baize
BARD	Mythological & Historical References	2023	Western	Google's BARD

Bianque	Mythological & Historical References	2023	China	BianQue: Balancing the Questioning and Suggestion Ability of Health LLMs with Multi-turn Health Conversations Polished by ChatGPT
BLOOM	Flora and Fauna Names	2023	Western	BigScience Large Open-science Open-access Multilingual Language Model
BloombergGPT	GPT + Domain-Specific Names	2023	Western	BloombergGPT
Caozhi	Mythological & Historical References	2023	China	{C}ao {Z}hi {L}arge {M}odel-{F}ocus on long-form text, multilingualism, and verticalization
Cerebras-GPT	GPT + Domain-Specific Names	2023	Western	Cebebras
Chameleon	Flora and Fauna Names	2023	Western	Chameleon
ChatGLM	Capability-driven names	2023	China	Chat General Language Model (GLM)
ChatGLM2	Derivation from original	2023	China	Chat General Language Model (GLM)
OpenBuddy	Capability-driven names	2023	China	OpenBuddy - Open Multilingual Chatbot for Everyone
Code Llama	Derivation from original	2023	Western	Code llama: Open foundation models for code
Dolly	Flora and Fauna Names	2023	Western	Free Dolly: Introducing the World's First Truly Open Instruction-Tuned LLM
Dromedary	Flora and Fauna Names	2023	Western	Dromedary
EduChat	GPT + Domain-Specific Names	2023	China	Educhat: A large-scale language model-based chatbot system for intelligent education
EmoGPT	GPT + Domain-Specific Names	2023	China	{E}mo{G}{P}{T} - {S}ecure {C}hat{G}{P}{T} {E}mail {W}riter for {G}mail
Falcon	Flora and Fauna Names	2023	Western	Falcon
ChatYuan	Scientific / Conceptual Names	2023	China	ChatYuan: Large Language Model for Dialogue in Chinese and English
Falcon-Instruct	Derivation from original	2023	Western	Falcon Instruct
Gemini	English Names	2023	Western	Gemini
Giraffe	Flora and Fauna Names	2023	Western	Giraffe: Adventures in expanding context lengths in llms
Hou Yi	Mythological & Historical References	2023	China	HouYi: An open-source large language model specially designed for renewable energy and carbon neutrality field
HuatuoGPT	GPT + Domain-Specific Names	2023	China	HuatuoGPT: A Medical LLM for Chinese Healthcare

Kosmos-1	Scientific / Conceptual Names	2023	Western	Kosmos
LLaMA	Flora and Fauna Names	2023	Western	Large Language Model Meta AI
LLaMA 2-chat	Derivation from original	2023	Western	See Llama
LLaMA-2	Derivation from original	2023	Western	See Llama
Mamba	Flora and Fauna Names	2023	Western	Mamba
MathGPT	GPT + Domain-Specific Names	2023	China	{M}ath{G}{P}{T} - {A}{I}{M}ath {S}olver - {M}ath {S}olver & {H}omework {H}elper
medGPT	GPT + Domain-Specific Names	2023	China	{H}ealthcare on autopilot - {M}edchat·ai
Mistral	Scientific / Conceptual Names	2023	Western	Mistral 7b
Nova	Scientific / Conceptual Names	2023	Western	Nova: Generative Language Models for Assembly Code with Hierarchical Attention and Contrastive Learning
OpenChat	Capability-driven names	2023	Western	OpenChat
Orca	Flora and Fauna Names	2023	Western	Orca: Progressive learning from complex explanation traces of gpt-4
Orca 2	Derivation from original	2023	Western	Orca 2
PaLM-E	Capability-driven names	2023	Western	PaLM-E
Phi-1	Scientific / Conceptual Names	2023	Western	Phi
Phi-2	Scientific / Conceptual Names	2023	Western	Phi
Phoenix	Mythological & Historical References	2023	China	Phoenix: Democratizing chatgpt across languages
Pythia	Mythological & Historical References	2023	Western	Pythia
RWKV	Capability-driven names	2023	Western	RWKV
SEEChat	Company branding names	2023	China	SEEChat: A Multimodal Conversational Model
StableLM	Company branding names	2023	Western	Stability Language Model
TigerBot	Flora and Fauna Names	2023	China	Tigerbot: An open multilingual multitask llm
Vicuna	Flora and Fauna Names	2023	Western	Vicuna

Wenxi Yiyan	Mythological & Historical References	2023	China	ERNIE 3.0 Titan
Wu Dao	Traditional Concepts	2023	China	Wu Dao
Xihu	Geographical names	2023	China	{X}ihu {X}incheng-{M}ultimodal large model empowering enterprises' intelligent transformation
Xuanyuan	Mythological & Historical References	2023	China	XuanYuan: A Large Chinese Financial Chat Model with Hundreds of Billions Parameters
Yuanchengxiang	Scientific / Conceptual Names	2023	China	YUAN 2.0: A Large Language Model with Localized Filtering-based Attention
Zhihaitu	Flora and Fauna Names	2023	China	{Z}hihai{T}u {C}hat-{Z}hihai{T}u {C}hat is an intelligent conversational assistant powered by a large language model under the {Z}hihu platform. — aibase.com
MiniGPT-4	GPT + Domain-Specific Names	2023	Western	MiniGPT-4: Enhancing Vision-language Understanding with Advanced Large Language Models
LLaVA	Capability-driven names	2023	Western	Large Language and Vision Assistant
VILA	Capability-driven names	2023	Western	Visual Language Models (VLMs)
BLIP-2	Derivation from original	2023	Western	Bootstrapping Language-Image Pre-training with Frozen Image Encoders and Large Language Models
LlongMa	Derivation from original	2023	Western	Long context Large Language Model Application
Arctic	Geographical names	2024	Western	Snowflake Arctic LLM
Athene	Mythological & Historical References	2024	Western	Athene
Baichuan 4	Derivation from original	2024	China	Baichuan: Open large-scale language models
CityGPT	GPT + Domain-Specific Names	2024	China	CityGPT: Empowering Urban Spatial Cognition of Large Language Models
Claude	Personal, historical, or scientific names	2024	Western	Claude
Einstein	Personal, historical, or scientific names	2024	Western	Einstein
CAIMIRA	Personal, historical, or scientific names	2024	Western	CAIMIRA
DBRX	Capability-driven names	2024	Western	DataBricks RX

deepseek	Company branding names	2024	China	Deepseek llm: Scaling open-source language models with longtermism
EXAONE 3.0	Scientific / Conceptual Names	2024	Western	EXAONE 3.0 by LG AI
Gemma	Mythological & Historical References	2024	Western	Gemma: Open Models Based on Gemini Research and Technology
GPT-4 Turbo	GPT + Domain-Specific Names	2024	Western	GPT-4
Granite	Scientific / Conceptual Names	2024	Western	Granite
Grok	Popular culture references	2024	Western	Grok
Hunyuan	Traditional Concepts	2024	China	Hunyuan-Large: An Open-Source MoE Model with 52 Billion Activated Parameters by Tencent
Hunyuan-DiT	Traditional Concepts	2024	China	Hunyuan-DiT: A Powerful Multi-Resolution Diffusion Transformer with Fine-Grained Chinese Understanding
Jiutian	Mythological & Historical References	2024	China	LION : Empowering Multimodal Large Language Model with Dual-Level Visual Knowledge
KOALA	Flora and Fauna Names	2024	Western	Key frame-conditioned long video-Large Language Model
Luminous 2	Scientific / Conceptual Names	2024	Western	Luminous
Mamba2	Derivation from original	2024	Western	Mamba2
Minerva	Mythological & Historical References	2024	Western	Minerva
Mixtral	Derivation from original	2024	Western	Mixture of Experts (MoE)
Nemotron	Popular culture references	2024	Western	Nemotron
o1	Company branding names	2024	Western	OpenAI's Q* or Strawberry
OLMo	Capability-driven names	2024	Western	Open Language Model
OLMoE	Derivation from original	2024	Western	Open Mixture-of-Experts Language Models
Shusheng	Traditional Concepts	2024	China	The Shusheng– Puyou 3.0 model (\url{https://www.aibase.com/news/14747})
TeleChat	Company branding names	2024	China	{T}ele{C}hat: An Open-source Billingual Large Language Model
Tianshu	Mythological & Historical References	2024	China	Tianshu: Towards Accurate Measuring, Modeling and Simulation of Deep Neural Networks

Tongyiqianwen	Traditional Concepts	2024	China	Tongyi Qianwen (or Qwen)
Zephyr	Scientific / Conceptual Names	2024	Western	Zephyr: Direct Distillation of {LM} Alignment
Baichuan Omni 1.5	Derivation from original	2025	China	Baichuan Omni 1.5
Claude 4.1	Derivation from original	2025	Western	Claude
Command	Capability-driven names	2025	Western	Command: optimized for conversational interaction and long context tasks
DeepSeek-V3.2-Exp	Derivation from original	2025	China	Deepseek llm: Scaling open-source language models with longtermism
Gemini 2.5 Pro	Derivation from original	2025	Western	Gemini
GLM-4.6	Derivation from original	2025	China	Generalized linear model
GPT-5	GPT + Domain-Specific Names	2025	Western	Generative Pre-trained Transformer
GPT-o3	GPT + Domain-Specific Names	2025	Western	Generative Pre-trained Transformer
GPT-o4-mini	GPT + Domain-Specific Names	2025	Western	Generative Pre-trained Transformer
Grok-5	Derivation from original	2025	Western	Generative Pre-trained Transformer
HunyuanWorld	Scientific / Conceptual Names	2025	China	HunyuanWorld
Jamba	Derivation from original	2025	Western	Jamba: Hybrid transformer-mamba language models
Kimi-K2	Scientific / Conceptual Names	2025	China	Kimi-K2
Llama 4	Derivation from original	2025	Western	Llama 4
Llama 4 Scout	Derivation from original	2025	Western	Llama 4 Scout
MAGMA	Geographical names	2025	Western	Multimodal Agentic Model
Qwen 3	Derivation from original	2025	China	Qwen
Sonus	Scientific / Conceptual Names	2025	Western	Sonus-1: A New Era of Large Language Models (LLMs)
BLIP3-o	Derivation from original	2025	Western	BLIP3-o: A Family of Fully Open Unified Multimodal Models-Architecture, Training and Dataset
Mixtral 2	Derivation from original	2025	Western	Mixture of Experts 2

Table 1: **A Cross-Cultural Census of Large Language Model Names** This table catalogs over 100 large language models, detailing their names, origins, release years, naming categories, and references. It highlights the contrasting naming conventions between Chinese and Western models, serving as the empirical foundation for analyzing cultural, functional, and strategic influences on LLM naming practices.

C.0.1 Fine-tuned or distilled versions of existing models:

1. DistilBERT – A smaller, faster, and lighter version of BERT, trained using knowledge distillation.
2. ALBERT – A lighter, optimized variant of BERT with parameter-sharing techniques.
3. DeBERTa – An improved BERT model with disentangled attention and enhanced positional encoding.
4. Flan-T5 – A fine-tuned version of T5, specifically trained on instruction-following tasks.
5. Flan-PaLM – A fine-tuned version of PaLM for better instruction tuning.
6. OPT-IML – A fine-tuned version of OPT, optimized for instruction-based tasks.
7. mT5 – A multilingual version of T5, designed for cross-lingual tasks.
8. UL2 – A fine-tuned version of T5 with a mixture of different pre-training objectives.
9. LLaMA 2-chat – A fine-tuned version of LLaMA 2, optimized for chatbot applications.
10. Hunyuan-DiT - Diffusion Transformer (DiT) (Hybrid vision-language model)
11. Baichuan - LLaMA-like (Chinese adaptation of LLaMA)
12. deepseek- LLaMA-inspired (Trained from scratch but similar framework)

C.0.2 Derived from other models but expanded:

1. Codex – Built on top of GPT-3, fine-tuned for code generation tasks (e.g., GitHub Copilot).
2. Claude – Likely influenced by GPT-4 and other transformer-based models but developed by Anthropic.
3. Gemini – Built by Google DeepMind, likely influenced by PaLM and T5.
4. ERNIE 3.0 – A more advanced version of ERNIE, a BERT-based model with knowledge integration.
5. GShard – A large-scale distributed training version of T5 and related transformer architectures.
6. GLaM – A mixture-of-experts model related to T5 and PaLM.
7. MT-NLG – A large-scale GPT-like model developed by Microsoft and NVIDIA.
8. Big BIRD – An optimized version of BERT for longer sequences.

D Annotation Procedure and Limitations

The annotation taxonomy was predefined before large-scale annotation and consisted of a fixed set of naming categories with explicit decision rules. These category definitions were then applied consistently across the full corpus. Two author-annotators independently categorized all 157 model names, and disagreements were resolved through discussion and reference to the predefined criteria.

Our annotation scheme is intentionally reception-oriented: rather than attempting to recover the private intent of model developers, it focuses on the public-facing semantic associations that names evoke once they enter discourse. This distinction is important because model names function socially as signals that shape interpretation, memory, and expectation. Ambiguous cases (e.g., *Gemma*) are therefore retained as meaningful boundary conditions that illustrate how names can support multiple plausible public interpretations simultaneously.

Agreement between the two primary annotators was high (150/157 exact matches, 95.5%, Cohen's $\kappa \approx 0.92$), suggesting that the taxonomy is stable and consistently applicable. To validate the annotation scheme beyond the author-annotators, two additional independent annotators applied the same predefined category definitions. These annotators agreed with the final inferred labels for 140/157 and 145/157 models, respectively, indicating that the categories generalize beyond the original annotation pair.

We acknowledge that this process does not substitute for a broader cross-cultural perception study. Future work could extend the analysis through large-scale surveys involving annotators and users from a wider range of linguistic and cultural backgrounds to evaluate how model names are interpreted globally.

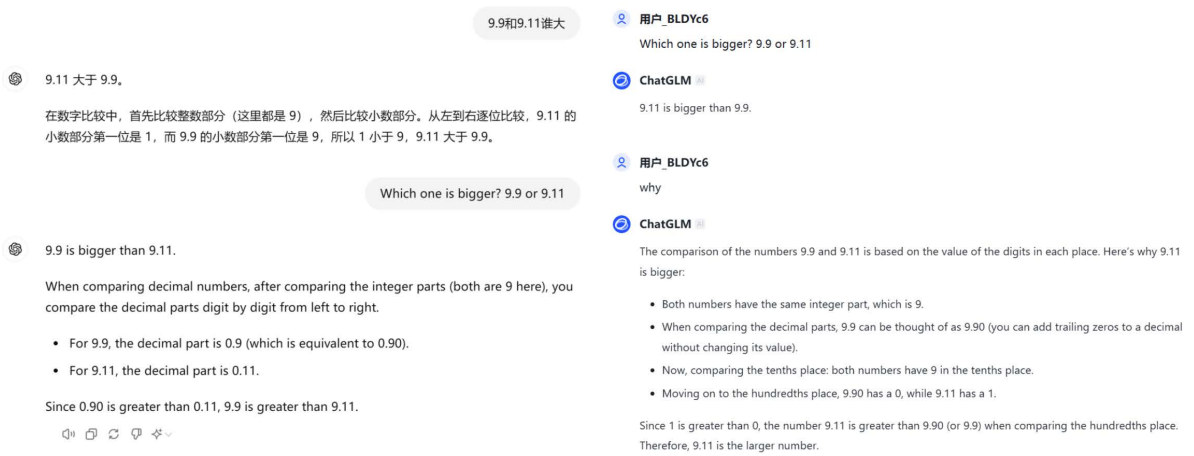

Figure 2: This figure shows how language models can generate incorrect answers while presenting them with seemingly logical reasoning. In the example, both GPT and ChatGLM incorrectly conclude that 9.11 is greater than 9.9. When asked in Chinese, GPT reasons step by step—comparing integer and decimal parts—but draws the wrong conclusion, while ChatGLM provides a similarly detailed yet flawed explanation. In English, however, GPT gives the correct answer. Because this example involves a simple comparison, it is easy to spot the error. But in more complex domains, users may not know the correct answer, making them more likely to trust the model's confident but mistaken reasoning.

D.1 Annotator Instructions

The following instructions were provided to annotators during the labeling process.

Task Overview You will be given the name of a large language model together with supporting context, such as the paper title, abstract, release announcement, or organization name. Your task is to assign the model to the single naming category that best explains the origin or motivation of the name.

The goal is not to classify the model architecture or technical capability, but rather to identify the naming strategy reflected in the model's public name.

Available Categories

1. **Muppet Character-Inspired Names**
   Names referencing Sesame Street or Muppet characters

2. **Capability-Driven Names**
   Names emphasizing functionality, architecture, scale, or technical capability.

3. **Popular Culture References**
   Names derived from movies, television, science fiction, games, or broader popular culture.

4. **Derivations from Existing Models**
   Names explicitly derived from earlier model names through prefixes, suffixes, version numbers, or modifications.

5. **Mythological & Historical References**
   Names based on mythological figures, legendary characters, philosophers, rulers, or historical individuals.

6. **Personal, Historical, or Scientific Names**
   Human given names or names associated with scientists, scholars, or recognizable individuals.

7. **GPT + Domain-Specific Names**
   Names combining "GPT" with a specialized application domain or topic.

8. **Flora and Fauna Names**
   Names inspired by animals, plants, or biological entities.

9. **Traditional Concepts**
   Names referencing philosophical, ethical, literary, or cultural concepts rather than specific individuals or places.

10. **Scientific / Conceptual Names**
    Names derived from scientific terminology, abstract concepts, natural phenomena, or metaphors.

11. **Geographical Names**
    Names based on geographic locations, landmarks, or regions.

12. **Company Branding Names**
    Names foregrounding organizational or corporate identity.

Annotation Guidelines

- Assign the category that best reflects the *primary naming motivation*.
- Use supporting documentation whenever possible.
- If a name could plausibly belong to multiple categories, choose the most salient or publicly emphasized interpretation.
- Prioritize explicit naming references over speculative interpretations.
- Version numbers or minor suffixes alone do not create new categories; they usually indicate derivation.
- If uncertain, annotators should leave a short justification note for later adjudication.

D.2 Limitations

The annotation process necessarily involved interpretive judgment. Some model names admit multiple plausible interpretations, particularly when acronyms simultaneously encode technical meanings and cultural references. In several cases, developers did not publicly explain the origin of a model name, requiring annotators to infer intent from contextual evidence such as documentation, interviews, or naming conventions within related model families.

The taxonomy itself also introduces limitations. Naming categories are not always mutually exclusive; for example, a model may simultaneously function as a corporate brand and a scientific metaphor. To improve consistency, annotators were instructed to assign the category representing the dominant or most publicly salient naming rationale.

Finally, the dataset is not exhaustive. The census focuses on widely discussed large language models and representative naming patterns rather than every released model worldwide. As naming conventions continue to evolve rapidly, future datasets may identify additional categories or reinterpret existing classifications.

Model name	Type	Year	Source	Published paper or technical report	Public availability
BERT	Muppet Character-Inspired Names	2018	Western	(Devlin et al., 2019b)	https://huggingface.co/docs/transformers/en/model_doc/bert
RoBERTA	Derivation from original	2018	Western	(Liu et al., 2019b)	https://huggingface.co/docs/transformers/en/model_doc/roberta
ELMO	Muppet Character-Inspired Names	2018	Western	(Peters et al., 2018b)	https://github.com/yuanxiaosc/ELMo
Trransformer-XL	Capability-driven names	2019	Western	(Dai et al., 2019)	https://huggingface.co/docs/transformers/v4.16.2/en/model_doc/transfo-xl
XLNet	Capability-driven names	2019	Western	(Yang et al., 2019)	https://huggingface.co/docs/transformers/en/model_doc/xlnet
ELECTRA	Mythological & Historical References	2019	Western	(Clark et al., 2020)	https://huggingface.co/docs/transformers/en/model_doc/electra
UniLM	Capability-driven names	2019	Western	(Dong et al., 2019)	https://github.com/microsoft/unilm
T5	Capability-driven names	2019	Western	(Raffel et al., 2020b)	https://huggingface.co/docs/transformers/en/model_doc/t5
DistilBERT	Derivation from original	2019	Western	(Sanh et al., 2019)	https://huggingface.co/docs/transformers/en/model_doc/distilbert
GPT	GPT + Domain-Specific Names	2019	Western	(Radford et al., 2018)	https://huggingface.co/docs/transformers/en/model_doc/openai-gpt
ERNIE	Muppet Character-Inspired Names	2019	China	(Zhang et al., 2019b)	https://huggingface.co/docs/transformers/en/model_doc/ernie
KERMIT	Muppet Character-Inspired Names	2019	Western	(Zanzotto et al., 2020)	https://github.com/ART-Group-it/KERMIT?tab=readme-ov-file

Grover	Muppet Character-Inspired Names	2019	Western	(Zellers et al., 2019b)	https://github.com/rowanz/grover
Rosita	Muppet Character-Inspired Names	2019	Western	(Mulcaire et al., 2019)	https://github.com/pmulcaire/rosita
Megatron-LM	Popular culture references	2019	Western	(Shoeybi et al., 2019)	https://github.com/NVIDIA/Megatron-LM
DeBERTa	Derivation from original	2020	Western	(He et al., 2021)	https://huggingface.co/docs/transformers/en/model_doc/deberta
mT5	Derivation from original	2020	Western	(Xue et al., 2021)	https://huggingface.co/docs/transformers/en/model_doc/mt5
GPT-Neo	Popular culture references	2020	Western	(Black et al., 2022a)	https://huggingface.co/docs/transformers/en/model_doc/gpt_neo
Reformer	Capability-driven names	2020	Western	(Kitaev et al., 2020)	https://huggingface.co/docs/transformers/en/model_doc/reformer
Pegasus	Mythological & Historical References	2020	Western	(Zhang et al., 2020a)	https://huggingface.co/docs/transformers/en/model_doc/pegasus
ALBERT	Derivation from original	2020	Western	(Lan et al., 2020)	https://huggingface.co/docs/transformers/en/model_doc/albert
DialoGPT	GPT + Domain-Specific Names	2020	Western	(Zhang et al., 2020b)	https://huggingface.co/docs/transformers/en/model_doc/dialogpt
Big BIRD	Muppet Character-Inspired Names	2020	Western	(Zaheer et al., 2020a)	https://huggingface.co/docs/transformers/en/model_doc/big_bird
MARGE	Muppet Character-Inspired Names	2020	Western	(Lewis et al., 2020a)	https://github.com/lucidrains/marge-pytorch
BART	Popular culture references	2020	Western	(Lewis et al., 2020b)	https://huggingface.co/docs/transformers/en/model_doc/bart

GShard	Capability-driven names	2021	Western	(Lepikhin et al., 2021)	https://github.com/lucidrains/mixture-of-experts
Codex	Capability-driven names	2021	Western	(Chen et al., 2021)	https://github.com/openai/codex
FLAN	Capability-driven names	2021	Western	(Wei et al., 2022)	https://github.com/google-research/FLAN
MT-NLG	Capability-driven names	2021	Western	(Smith et al., 2022)	https://github.com/NVIDIA/Megatron-LM
GPT-J	GPT + Domain-Specific Names	2021	Western	(Wang and Komatsuzaki, 2021)	https://huggingface.co/docs/transformers/en/model_doc/gptj
GPT-NeoX-20B	Derivation from original	2021	Western	(Black et al., 2022b)	https://huggingface.co/EleutherAI/gpt-neox-20b
CPM	Capability-driven names	2021	China	(Zhang et al., 2021a)	https://huggingface.co/docs/transformers/en/model_doc/cpm
CPM-2	Derivation from original	2021	China	(Zhang et al., 2021a)	https://github.com/TsinghuaAI/CPM
GLaM	Capability-driven names	2021	Western	(Du et al., 2022a)	https://huggingface.co/FinchResearch/GLaM
WebGPT	GPT + Domain-Specific Names	2021	Western	(Nakano et al., 2021)	https://github.com/0hq/WebGPT
EVA	Capability-driven names	2021	China	(Zhou et al., 2021)	https://huggingface.co/BAAI/EVA
Pangu	Mythological & Historical References	2021	China	(Zeng et al., 2021)	https://github.com/huawei-noah/Pretrained-Language-Model/tree/master/PanGu-\alpha
Mengzi	Mythological & Historical References	2021	China	(Zhang et al., 2021b)	https://www.aibase.com/tool/29143(493million)
Jurassic-1	Scientific / Conceptual Names	2021	Western	(Lieber et al., 2021)	https://www.ai21.com/blog/ai21-studio-use-cases/
WenLan	Traditional Concepts	2021	China	(Huo et al., 2021)	https://github.com/BAAI-WuDao/BriVl
Flan-T5	Capability-driven names	2022	Western	(Chung et al., 2024)	https://huggingface.co/docs/transformers/en/model_doc/flan-t5
Flan-PaLM	Capability-driven names	2022	Western	(Chung et al., 2024)	https://huggingface.co/google/flan-t5-xxl

OPT-IML	Capability-driven names	2022	Western	(Iyer et al., 2023)	https://huggingface.co/facebook/opt-iml-1.3b
UL2	Capability-driven names	2022	Western	(Tay et al., 2023)	https://huggingface.co/facebook/opt-iml-1.3b
OPT	Capability-driven names	2022	Western	(Zhang et al., 2022)	https://huggingface.co/docs/transformers/en/model_doc/opt
T0	Capability-driven names	2022	Western	(Sanh et al., 2022)	https://github.com/bigscience-workshop/t-zero
Retro	Capability-driven names	2022	Western	(Borgeaud et al., 2022)	https://github.com/sunnynexus/RetroLLM
LaMDA	Capability-driven names	2022	Western	(Thoppilan et al., 2022)	https://github.com/conceptofmind/LaMDA-rlhf-pytorch
GLM-130B	Capability-driven names	2022	Western	(Du et al., 2022b)	https://github.com/THUDM/GLM-130B
Chinchilla	Flora and Fauna Names	2022	Western	(Hoffmann et al., 2022)	https://github.com/kyo-takano/chinchilla
Sparrow	Flora and Fauna Names	2022	Western	(Glaese et al., 2022)	https://github.com/katanaml/sparrow
BioGPT	GPT + Domain-Specific Names	2022	Western	(Luo et al., 2022)	https://huggingface.co/docs/transformers/en/model_doc/biogpt
InstructGPT	GPT + Domain-Specific Names	2022	Western	(Ouyang et al., 2022)	https://github.com/openai/following-instructions-human-feedback
Gopher	Muppet Character-Inspired Names	2022	Western	(Rae et al., 2021)	https://llmmodels.org/tools/gopher/
Bagualu	Mythological & Historical References	2022	China	(Ma et al., 2022)	https://hackmd.io/@enochliang/SkfCeQ07n
Galactica	Scientific / Conceptual Names	2022	Western	(Taylor et al., 2022)	https://huggingface.co/facebook/galactica-1.3b
ChatGLM	Derivation from original	2023	China	(Zeng et al., 2024)	https://github.com/zai-org/ChatGLM-6B
PaLM-E	Capability-driven names	2023	Western	(Driess et al., 2023)	https://github.com/kyegomez/PALM-E
RWKV	Capability-driven names	2023	Western	(Peng et al., 2023)	https://github.com/rwkv

Phoenix	Mythological & Historical References	2023	China	(Chen et al., 2023c)	https://github.com/Arize-ai/phoenix
TigerBot	Flora and Fauna Names	2023	China	(Chen et al., 2023a)	https://huggingface.co/TigerResearch
Code Llama	Derivation from original	2023	Western	(Roziere et al., 2023)	https://github.com/meta-llama/codellama
LLaMA-2	Derivation from original	2023	Western	(Touvron et al., 2023b)	https://huggingface.co/docs/transformers/en/model_doc/llama2
LLaMA 2-chat	Derivation from original	2023	Western	(Touvron et al., 2023b)	https://huggingface.co/docs/transformers/en/model_doc/llama2
Gemini	English Names	2023	Western	(Team et al., 2023)	https://huggingface.co/describeai/gemini
Giraffe	Flora and Fauna Names	2023	Western	(Pal et al., 2023)	https://github.com/autonomousvision/giraffe
LLaMA	Flora and Fauna Names	2023	Western	(Touvron et al., 2023a)	https://github.com/meta-llama/llama
BLOOM	Flora and Fauna Names	2023	Western	(Le Scao et al., 2023)	https://huggingface.co/bigscience/bloom
Dromedary	Flora and Fauna Names	2023	Western	(Sun et al., 2023)	https://github.com/IBM/Dromedary
Mamba	Flora and Fauna Names	2023	Western	(Gu and Dao, 2023)	https://huggingface.co/docs/transformers/main/en/model_doc/mamba
Vicuna	Flora and Fauna Names	2023	Western	(Chiang et al., 2023)	https://huggingface.co/lmsys/vicuna-7b-v1.5
StableLM	Company Names	2023	Western	(Bellagente et al., 2024)	https://huggingface.co/docs/transformers/en/model_doc/stablelm
Orca	Flora and Fauna Names	2023	Western	(Mukherjee et al., 2023)	https://huggingface.co/microsoft/Orca-2-13b
Orca 2	Derivation from original	2023	Western	(Mitra et al., 2023)	https://huggingface.co/microsoft/Orca-2-13b
Phi-1	Scientific / Conceptual Names	2023	Western	(Gunasekar et al., 2023)	https://huggingface.co/microsoft/phi-1
OpenChat	Capability-driven names	2023	Western	(Wang et al., 2024a)	https://github.com/imoneoi/openchat

Falcon	Flora and Fauna Names	2023	Western	(Almazrouei et al., 2023)	https://huggingface.co/docs/transformers/main/en/model_doc/falcon
Chameleon	Flora and Fauna Names	2023	Western	(Lu et al., 2023)	https://huggingface.co/docs/transformers/en/model_doc/chameleon
Alpaca	Flora and Fauna Names	2023	Western	(Taori et al., 2023)	https://github.com/tatsu-lab/stanford_alpaca
Dolly	Flora and Fauna Names	2023	Western	(Borisov and Mikhaylovskiy, 2023)	https://huggingface.co/databricks/dolly-v2-12b
Cerebras-GPT	GPT + Domain-Specific Names	2023	Western	(Dey et al., 2023)	https://huggingface.co/cerebras/Cerebras-GPT-13B
BloombergGPT	GPT + Domain-Specific Names	2023	Western	(Wu et al., 2023b)	https://belitsoft.com/bloomberggpt
MathGPT	GPT + Domain-Specific Names	2023	China	(Scarlatos and Lan, 2023)	https://github.com/umass-ml4ed/mathGPT
EduChat	GPT + Domain-Specific Names	2023	China	(Dan et al., 2023)	https://edu.chat/
medGPT	GPT + Domain-Specific Names	2023	China	(Kraljevic et al., 2021)	https://www.medgpt.app/
HuatuoGPT	GPT + Domain-Specific Names	2023	China	(Zhang et al., 2023)	https://github.com/FreedomIntelligence/HuatuoGPT
Atlas	Mythological & Historical References	2023	Western	(Izacard et al., 2023)	https://github.com/facebookresearch/atlas
Pythia	Mythological & Historical References	2023	Western	(Biderman et al., 2023)	https://github.com/EleutherAI/pythia
BARD	Mythological & Historical References	2023	Western	(Ram and Verma, 2023)	https://github.com/Illia-the-coder/BardAPI-LLM
Baize	Mythological & Historical References	2023	Western	(Xu et al., 2023)	https://github.com/project-baize/baize-chatbot

Wenxi Yiyan	Mythological & Historical References	2023	China	(Baidu, 2025)	https://yiyan.baidu.com/(1.5 billion)[6]
Xuanyuan	Mythological & Historical References	2023	China	(Zhang and Yang, 2023)	https://github.com/Duxiaoman-DI/XuanYuan
Caozhi	Mythological & Historical References	2023	China	(Alibaba, 2025)	https://www.aibase.com/tool/28994(22500)
Bianque	Mythological & Historical References	2023	China	(Chen et al., 2023b)	https://github.com/scutcyr/BianQue
Mistral	Scientific / Conceptual Names	2023	Western	(Jiang et al., 2023a)	https://huggingface.co/mistralai/Mistral-7B-v0.1
Kosmos-1	Scientific / Conceptual Names	2023	Western	(Huang et al., 2023)	https://gpt3demo.com/apps/microsoft-kosmos-1
Phi-2	Scientific / Conceptual Names	2023	Western	(Javaheripi et al., 2023)	https://huggingface.co/microsoft/phi-2
Nova	Scientific / Conceptual Names	2023	Western	(Jiang et al., 2025)	https://huggingface.co/lt-asset/nova-1.3b-bcr
Yuanchengxiang	Scientific / Conceptual Names	2023	China	(Wu et al., 2023a)	https://github.com/Shawn-IEITSystems/Yuan-1.0
Wu Dao	Traditional Concepts	2023	China	(BAAI, 2021)	https://github.com/BAAI-WuDao
Hou Yi	Mythological & Historical References	2023	China	(Bai et al., 2023)	https://github.com/LLMSecurity/HouYi
Baichuan	Traditional Concepts	2023	China	(Yang et al., 2023)	https://huggingface.co/baichuan-inc/Baichuan-7B
MiniGPT-4	GPT + Domain-Specific Names	2023	Western	(Ataallah et al., 2024)	https://minigpt-4.github.io/
BLIP-2	Capability-driven names	2023	Western	(Li et al., 2023)	https://github.com/salesforce/BLIP
AceGPT	GPT + Domain-Specific Names	2023	China	(Huang et al., 2024)	https://huggingface.co/FreedomIntelligence/AceGPT-13B
Aquila	Flora and Fauna Names	2023	China	(Lu et al., 2024)	https://huggingface.co/BAAI/Aquila-7B

[6] https://www.aibase.com/news/13702

Mixtral	Derivation from original	2024	Western	(Jiang et al., 2024)	https://huggingface.co/mistralai/Mixtral-8x7B-v0.1
OLMo	Capability-driven names	2024	Western	(Groeneveld et al., 2024)	https://huggingface.co/docs/transformers/en/model_doc/olmo
o1	Company Names	2024	Western	(Zhong et al., 2024)	https://openai.com/o1/
TeleChat	Company Names	2024	China	(Wang et al., 2024b)	https://github.com/Tele-AI/Telechat
Chimed-gpt	GPT + Domain-Specific Names	2024	China	(Tian et al., 2024)	https://github.com/synlp/ChiMed-GPT
deepseek	Company Names	2024	China	(Bi et al., 2024)	https://huggingface.co/deepseek-ai
OLMoE	Derivation from original	2024	Western	(Muennighoff et al., 2024)	https://huggingface.co/docs/transformers/en/model_doc/olmoe
Mamba2	Derivation from original	2024	Western	(Dao and Gu, 2024)	https://huggingface.co/docs/transformers/en/model_doc/mamba2
KOALA	Flora and Fauna Names	2024	Western	(Zhang et al., 2024)	https://github.com/rxtan2/Koala-video-llm
CityGPT	GPT + Domain-Specific Names	2024	China	(Feng et al., 2024)	https://github.com/tsinghua-fib-lab/citygpt
Claude	Popular culture references	2024	Western	(Enis and Hopkins, 2024)	https://github.com/simonw/llm-claude-3
Nemotron	Popular culture references	2024	Western	(Parmar et al., 2024)	https://huggingface.co/docs/transformers/en/model_doc/nemotron
EXAONE 3.0	Scientific / Conceptual Names	2024	Western	(Research et al., 2024)	https://github.com/LG-AI-EXAONE/EXAONE-3.0
Minerva	Mythological & Historical References	2024	Western	(Orlando et al., 2024)	https://huggingface.co/collections/sapienzanlp/minerva-llms
DBRX	Capability-driven names	2024	Western	(Gupta and Yip, 2024)	https://huggingface.co/docs/transformers/en/model_doc/dbrx

Zephyr	Scientific / Conceptual Names	2024	Western	(Tunstall et al., 2024)	https://huggingface.co/HuggingFaceH4/zephyr-7b-beta
Gemma	Mythological & Historical References	2024	Western	(Mesnard et al., 2024)	https://huggingface.co/docs/transformers/en/model_doc/gemma
Granite	Scientific / Conceptual Names	2024	Western	(Mishra et al., 2024)	https://github.com/ibm-granite/granite-code-models
Tongyiqianwen	Traditional Concepts	2024	China	(Yang et al., 2024)	https://huggingface.co/Qwen
Hunyuan	Traditional Concepts	2024	China	(Sun et al., 2024)	https://huggingface.co/tencent/Tencent-Hunyuan-Large
Grok	Popular culture references	2024	Western	N/A	https://huggingface.co/xai-org/grok-1 (35.1 million[7])
Shusheng	Traditional Concepts	2024	China	(Cai et al., 2024)	https://github.com/vealocia
Jiutian	Mythological & Historical References	2024	China	(Chen et al., 2024)	https://github.com/JiuTian-VL/JiuTian-LION
Hunyuan-DiT	Traditional Concepts	2024	China	(Li et al., 2024)	https://github.com/Tencent/HunyuanDiT
Arctic	Geographical names	2024	Western	(Borchmann et al., 2025)	https://github.com/Snowflake-Labs/snowflake-arctic
Command	Capability-driven names	2025	Western	(Cohere et al., 2025)	https://cohere.com/command
Grok 5	Derivation from original	2025	Western	N/A	https://huggingface.co/xai-org/grok-1(35.1 million[8])
Jamba	Derivation from original	2025	Western	(Lenz et al., 2025)	https://huggingface.co/ai21labs/Jamba-v0.1
Sonus	Scientific / Conceptual Names	2025	Western	(Chat, 2025)	https://sonus.ai/
GPT-5	GPT + Domain-Specific Names	2025	Western	(Wang et al., 2025)	https://openai.com/gpt-5/
MAGMA	Geographical names	2025	Western	(Yang et al., 2025b)	https://github.com/microsoft/Magma
Claude 4.1	Derivation from original	2025	Western	(Sorensen)	https://claude.ai/

[7] https://explodingtopics.com/blog/grok-users
[8] https://explodingtopics.com/blog/grok-users

Qwen 3	Derivation from original	2025	China	(Yang et al., 2025a)	https://huggingface.co/collections/Qwen/qwen3
GPT-o4-mini	GPT + Domain-Specific Names	2025	Western	(OpenAI et al., 2024)	https://platform.openai.com/docs/models/o4-mini
GPT-o3	GPT + Domain-Specific Names	2025	Western	(Kataoka et al., 2025)	https://platform.openai.com/docs/models/o3
Llama 4 Scout	Derivation from original	2025	Western	(Adcock et al., 2026)	https://huggingface.co/meta-llama/Llama-4-Scout-17B-16E
Gemini 2.5 Pro	Derivation from original	2025	Western	(Comanici et al., 2025)	https://deepmind.google/models/gemini/pro/
DeepSeek-V3.2-Exp	Derivation from original	2025	China	(Liu et al., 2025)	https://huggingface.co/deepseek-ai/DeepSeek-V3.2-Exp
Llama 4	Derivation from original	2025	Western	(Abdullah et al., 2025)	https://huggingface.co/collections/meta-llama/llama-4
Kimi-K2	Scientific / Conceptual Names	2025	China	(Bai et al., 2025)	https://github.com/MoonshotAI/Kimi-K2
HunyuanWorld	Scientific / Conceptual Names	2025	China	(Team et al., 2025)	https://github.com/Tencent-Hunyuan/HunyuanWorld-1.0
GLM-4.6	Derivation from original	2025	China	(Zeng et al., 2025)	https://huggingface.co/zai-org/GLM-4.6
BLIP3-o	Derivation from original	2025	Western	(Chen et al., 2025)	https://github.com/JiuhaiChen/BLIP3o

Table 2: This table presents the list of models that met at least two of the three key inclusion criteria used in the study. It serves as the final dataset for deeper cultural and naming convention analysis across Chinese and Western models.