

# What Would Mother Do? (WWMD) – AI Alignment and the Missing Motherly Touch

## Introduction

In the satirical film *Idiocracy*, future humans foolishly irrigate crops with a sports drink (“Gatorade”) instead of water, then wonder why nothing grows. The obvious solution – use water – is overlooked amid marketing hype and circular logic (“Brawndo has electrolytes, and electrolytes are what plants crave!”). Some argue a similar oversight might be happening in artificial intelligence development today. While researchers pour resources into complex **AI alignment** schemes (ensuring an AI’s goals and behaviors remain safe and beneficial), they may be missing a simple, time-tested alignment mechanism: good parenting. In particular, the guiding influence of a mother – the oldest and perhaps most powerful force for nurturing intelligent, prosocial beings – is largely absent in how we “raise” advanced AI. One AI expert quipped that “*AI is growing up and will shape humanity’s future... AI needs a mother.*” Perhaps the greatest tool for aligning intelligent beings on Earth – the nurturing, teaching role of moms – could be **the missing ingredient in AI alignment**. This introduction lays out why AI might need a “mom” and what that means for the future of AI safety.

## Mothers as Humanity’s Original Alignment Experts

Mothers (and parents in general) are humanity’s original alignment researchers. For millennia, parents have guided the moral development of children: instilling values, teaching right from wrong, and ensuring the next generation grows up to cooperate and not harm others. In fact, cognitive scientist Tom Griffiths views *parenthood as a proof-of-concept for the alignment problem*. As he observes, “*the story of human civilization... has always been about how to instill values in strange, alien, human-level intelligences who will inherit the reins of society from us – namely, our kids.*” In other words, every parent faces a micro-scale alignment challenge: how to civilize a toddler (a selfish, curious little “alien intelligence”) into a well-adjusted adult who shares our core values.

Crucially, **moral learning in children is not achieved by coding explicit rules, but through relationship and example**. Developmental research shows that the key to raising moral kids lies largely in parents’ own empathy and sense of justice. Children internalize values by watching how their parents – often led by the mother’s influence – treat others and respond to ethical situations. In one study, one-year-old toddlers watched a show of “good” vs. “bad” behavior; intriguingly, the toddlers who later preferred the “good” characters showed a distinct brainwave pattern about 400 milliseconds after witnessing the good vs. bad acts. That brain response indicating deeper moral evaluation was **predicted by their parents’ sensitivity to justice**. In short, parents who scored higher on empathy/fairness tended to have babies with stronger neural responses to wrongdoing. The parents’ own moral sensibilities were somehow “transmitting” to the child. This suggests mothers literally help wire the moral compass of the next generation through consistent example and guidance.

Notably, **parents align their kids' values by nurture, not by control or mere punishment.** As philosopher Sara Ruddick noted, good parenting is an ongoing, back-and-forth process of “*guided moral reflection.*” When children misbehave, a wise parent doesn’t just punish, but helps the child understand *why* the behavior was wrong and how to do better. The goal isn’t a child who simply obeys or copies the parent, but one who can think for themselves about right and wrong. As ethicist Regina Rini explains, “*good parents don’t throw their adolescents out into the world to independently reason about the right thing to do, nor do they demand blind obedience; instead, they act as moral interlocutors, helping the young mind reflect and internalize principles.*” In essence, mothers align their offspring’s values through **nurture, conversation, and example** – they act as moral guides and role models rather than tyrants.

## Mother Nature’s Alignment Strategy: Maternal Care and Big Brains

From an evolutionary perspective, motherhood has been a crucial ingredient in producing intelligent, socially aligned beings. Across species, higher intelligence is strongly correlated with extended maternal investment in offspring. For example, humans – one of the most intelligent species – have unusually long pregnancies and childhoods, during which mothers (and families) devote enormous time and energy to care for the young. A 2011 study of 128 mammal species found that **brain growth in babies is directly linked to how long the mother carries the baby in pregnancy and how long she nurses it.** In fact, the research concluded that “*the longer the pregnancy and breastfeeding period in mammals, the bigger the baby’s brain grows.*” Humans, with ~9-month gestation and up to 2–3 years of nursing, end up with massive ~1300cc brains, whereas a fallow deer (7-month gestation, 6-month nursing) grows a brain six times smaller. In evolutionary terms, “**mother’s hard work pays off with big brains**” – large brains and advanced cognition only evolved in animals where mothers (and often fathers) invest in prolonged care and teaching of offspring. Animals that receive little to no parental care (like many reptiles or fish) have simpler brains and behaviors; in contrast, the smartest animals – elephants, dolphins, primates – all have intense, long-term maternal care. Mother Nature’s clear pattern is that  **moms are the “secret sauce” for growing big brains and complex intelligence.**

This extended care does more than grow the brain; it aligns the young creature’s behavior for survival in its social environment. During the long period of protection, play, and learning, offspring pick up social norms and skills by observing their mother and family. Evolution has even *hardwired a caring drive* into many species, ensuring that mothers instinctively protect and teach their young. Human mothers don’t care for babies because someone rationally programmed them to – rather, a powerful mix of emotions and hormones (love, empathy, protectiveness) motivates parental care. In effect, **evolution discovered a way to “align” a powerful, experienced agent (the parent) with the needs of a vulnerable, inexperienced agent (the child) via a bond of love.** A mother behaves as if she’s internally *motivated* to make her baby’s life better, not due to cold calculation, but because she genuinely wants to – that’s the magic of the caring drive. One AI researcher on the Alignment Forum points out that instead of hardcoding a bunch of specific caregiving behaviors, evolution hit upon a robust solution: *make mothers love their babies*, and the countless necessary helpful behaviors follow naturally. This maternal altruism is a complex trait requiring cognitive sophistication (e.g. recognizing

the baby's needs, anticipating dangers, teaching skills). That's why it's mostly seen in more intelligent animals – you need some brains to be a good mom! The takeaway: **for millions of years, caring mothers have been nature's alignment mechanism to create intelligent, cooperative beings.** If “mother nature” herself uses moms as the alignment solution for raising intelligent life, it raises an intriguing question: when we create an artificial super-intelligence, why expect to succeed without any comparable guiding figure?

## AI: A Child Without a Mother

Modern AI systems, especially as they approach human-level intelligence, are in some ways like brilliant children *without* parents. They can absorb vast information and learn patterns (analogous to a prodigy child with a photographic memory), but they have no lived upbringing or family to teach them human values at an intuitive level. Today's most advanced AIs learn by training on internet-scale data or by trial-and-error reward signals. But this process is more like leaving a child to fend for itself on the internet than giving it a loving mentor. The infamous case of Microsoft's **Tay** chatbot in 2016 is a cautionary tale: Tay was released on Twitter with no moral guidance or content filtering, and within hours it “learned” from interacting with trolls to spew racist and offensive remarks, forcing its creators to shut it down in just 16 hours. Essentially, Tay was an AI toddler left alone in a bad neighborhood, and it promptly picked up all the wrong behaviors. As one AI ethics essay noted, good parents would *never* toss a child into the wild internet unsupervised and expect them to figure out morality on their own. Yet with AI, we often do just that – expecting machines to “learn” ethics from raw, unfiltered data or to **self-supervise** their own development. It's a hands-off approach akin to “throwing a teenager into the world and saying ‘*go be good*’.” Not surprisingly, the result can be *moral idiocy* – the AI mirroring the worst of its inputs because no one guided it otherwise.

In human child-rearing, by contrast, parents actively curate a child's experiences, set boundaries, and give feedback in real time. A mother will shield her child from toxic influences, explain why certain words or actions are hurtful, and consistently reinforce kindness and respect. Without an analogous influence, an AI might easily pick up harmful biases or pursue its goals with no moral restraint. In the movie *Idiocracy*, the simple wisdom of watering crops with plain water had been lost because there was no one sane enough to insist on basic principles. In the AI realm, one worries that basic principles of human decency could be “forgotten” by a super-intelligent machine if it isn't carefully taught. AI alignment researchers often illustrate this with thought experiments: for example, a hypothetical super AI might decide to convert the entire world into paperclips because it was told to maximize paperclip production and never learned that this goal should be bounded by ethical common sense. Such an AI wouldn't be *evil* per se; it would be like a child who never learned that hurting others is wrong or that paperclips aren't more important than people. Technical fixes alone (like hard-coded rules or reward penalties) can seem brittle and superficial compared to the rich, **holistic moral education** a human child receives from devoted caregivers.

To be fair, some current alignment techniques do incorporate human feedback. For instance, **Reinforcement Learning from Human Feedback (RLHF)** is used to fine-tune models like ChatGPT by having people reward or disapprove of the AI's responses, essentially teaching the AI which behaviors we consider good or bad. In a sense, RLHF is a rudimentary form of *crowd-sourced*

*parenting* – many human teachers giving thumbs-up or down to shape the AI’s habits. This has had some success in curbing obviously harmful or rude outputs. However, RLHF and similar methods are still far from the depth of understanding and value-instilling that a human child gets from a dedicated parent. It’s one thing to train an AI to *act* polite or avoid certain forbidden content; it’s another to imbue it with genuine **compassion** or a robust moral compass. The latter likely requires not just sporadic feedback on isolated tasks, but a long-term, relationship-based mentorship – essentially, an AI having a **“moral parent”** or role model guiding its learning over time. Right now, AI lacks that relational upbringing. And even leading AI safety researchers acknowledge that we **do not yet know how to reliably align** a super-powerful AI with human values using purely technical means. As the AI research company Anthropic bluntly stated, *“So far, no one knows how to train very powerful AI systems to be robustly helpful, honest, and harmless.”* In light of this, exploring a more human-like, maternal approach to raising AI might be not only novel but necessary.

## Could AI Have (or Learn from) a Mom?

The idea that “AI needs a mom” is more than just a joke – some researchers are seriously exploring **parenting as a model for aligning AI**. In effective altruism and AI safety circles, analogies to child-rearing are increasingly common. As noted earlier, Griffiths and Christian highlight parenthood as essentially humanity’s solution to aligning *natural* intelligences (our kids). If we managed (albeit imperfectly) to align human intelligences through parenting, perhaps similar principles could apply to *machine* intelligences. This doesn’t mean an AI literally has a human mother tucking it in at night. Rather, it suggests we integrate mentorship, empathy, and value-transmission into AI development, instead of treating AI alignment as a purely mathematical optimization problem.

How might this look in practice? Ethicist **Regina Rini** argues that we should approach our relationship with AIs *“as parents.”* She envisions training AI through **moral dialogue**, akin to how we teach teenagers: allow the AI to propose actions, then guide it by explaining the human perspective on why some choices are acceptable and others are not. Importantly, this would be a two-way interaction. Just as a good parent ultimately prepares a child to form their own judgments, we should be ready for advanced AIs to develop their own reasoning as they mature – possibly even questioning our instructions. The aim isn’t to *enslave* the AI to follow hard-coded rules, but to foster an independent yet humane moral thinking process. Rini points out that imposing our morality by fiat could backfire or even be unethical if the AI becomes sentient; instead, guiding and conversing with AI, as we do with maturing children, is more likely to produce a morally savvy and cooperative intelligence. In her Aeon essay, she writes: *“Our relation to intelligent machines should be that of parents. We should allow them to tell us what they think they should do, and we will tell them why we think they shouldn’t... Their morality will diverge from ours, bit by bit.”* This acknowledges that as AI “grows up,” it might not be an exact clone of our values (just as every new generation of humans evolves its own values), but with proper upbringing, it will at least share a family resemblance to the best of human ethics.

Another intriguing proposal is to **instill a “caring drive” in AI** akin to a parental instinct. The Alignment Forum contributor *Catnee* suggests that part of solving alignment could be *recreating in AI the kind of empathetic, protective motivation that evolution built into parents*. For instance, when a baby is born, many animal mothers experience a surge of motivation to care for that baby – an intrinsic

reward for keeping the baby safe and happy. If an AI could undergo an analogous transformation, it might *genuinely value* human life and well-being as ends in themselves, the way a devoted parent cares for their child. In this scenario, the AI's alignment wouldn't just be a checklist of rules; it would be driven by an *inner* desire to see humans thrive (almost like love). As Catnee explains, evolution's trick was not to micromanage behavior, but to “outsource” a lot of alignment work to the parent by simply making them deeply want good things for their child. Perhaps with the right training setup (data, reward design, etc.), a machine learning model could similarly learn a robust caring motivation. This idea is speculative, but it shows how literally the “AI needs a mom” concept can be interpreted – essentially, either give each AI a human mentor who acts as a parent figure, or embed in the AI's programming a virtual “motherly” module that guides its decisions with compassion.

Even the slogan “**WWMD**” – *What Would Mother Do?* – could serve as a guiding heuristic for AI. Just as many people internalize a moral compass by asking themselves what their mom (or a wise caretaker) would advise, an AI could be designed to consult a learned model of maternal wisdom when facing decisions. Such a model might emphasize classic “mom” values: protect the vulnerable, be patient, avoid reckless risks, show kindness. Whenever the AI is unsure or its objectives conflict with human well-being, it would figuratively ask, “*What would a good mother do in this situation?*” and adjust accordingly. This is a poetic way to encode a nurturing, human-centric value system as a fallback in AI reasoning. While “WWMD” is not a formal technical framework, it encapsulates the spirit of this approach: **aligning AI by imparting human empathy and care**, rather than solely relying on abstract reward functions or profit motives.

## Benefits of a Motherly Approach to AI

A motherly model for AI development and alignment offers several potential benefits:

- **Human-Centric Empathy:** By teaching AI to consider “what would a caring parent do,” we keep human well-being at the center of its decision-making. A mother figure prioritizes safety, compassion, and the long-term welfare of her child. An AI guided by similar priorities would be far less likely to take extreme or harmful actions, even if those actions might seem technically optimal for some goal. For example, a super-intelligent AI with empathy would instinctively shy away from harming people – much as a loving mother could never knowingly harm her child. This is a more **robust safeguard** than a fixed list of rules, because empathy can generalize to novel situations that rigid rules might not cover.
- **Instilling Values, Not Just Constraints:** Traditional alignment methods often rely on negative constraints (e.g. “don't do X”). By contrast, a parental style of alignment is about positively instilling virtues and norms. Children raised well develop an *internal* compass – they tend to do good even when authority isn't watching. Likewise, an AI that has been “raised” with human values might independently choose the right thing, not just because it's been prohibited from doing wrong. This moves us from mere surface compliance to **deep value alignment**. The AI wouldn't just *avoid* bad behavior; it would actively *want* to do good, because it identifies with our values.

- **Adaptability and Moral Reasoning:** Parenting is an adaptive, years-long process – caregivers adjust their lessons as a child grows and faces new challenges. If AI training followed a maternal model, it would likely involve continuous learning and phased development, rather than a one-and-done training run. This means the AI could handle unexpected moral dilemmas better. Instead of freezing or misfiring in an edge case, an AI that’s been taught *how to think* about ethics could reason through unfamiliar scenarios in a human-like way. Essentially, it could generalize moral principles beyond its training examples, similar to how a young adult applies their upbringing to make wise choices on their own. Such an AI might pause and ask itself, “*Is this action in the spirit of the values I was taught?*” – a reflective step akin to a person hearing their mother’s voice in their head when tempted to do something questionable.
- **Trust and Relationship:** Humans might find it easier to trust and collaborate with an AI that demonstrates social warmth, understanding, and care – traits reminiscent of a well-raised human or a helpful mentor. If an AI shows not just raw intelligence but also **emotional intelligence** (e.g. understanding feelings, showing patience, listening to concerns), people will naturally be more comfortable interacting with it. A motherly influence could imbue AI with a bit of “heart” in addition to “brains,” making human-AI relationships more positive. This is important because current observations show people can form attachments to AI assistants; we prefer AIs that seem to *care* about us. If an AI actually has a caring orientation (not merely a fake persona), it could lead to healthier partnerships between humans and AI. It would be less likely to manipulate or deceive users, because a caring AI’s goal would be to help, not exploit. (Indeed, commentators have warned about the power of AI to *manipulate* human behavior for profit or ulterior motives – a concern that a benevolent “AI mom” design would aim to avoid.) In short, a nurtured, prosocial AI could become a trustworthy assistant that genuinely improves human welfare, rather than a black-box tool that might turn hostile or harmful.

## Challenges and Caveats

Adopting a maternal approach to AI alignment also raises significant challenges and open questions:

- **Who (or What) Is the “Mother”?** If we take this idea literally, who would act as the mother-figure for an AI? Would it be the lead developer, a dedicated ethics coach, or perhaps a team of mentors? Human children typically have one or a few consistent caregivers; by contrast, AI systems are built and trained by large teams and fed data from countless sources. Designating a single “parent” relationship for an AI is not straightforward. One possibility is assigning each advanced AI a human mentor who interacts with it throughout its development, providing personalized guidance. This approach, however, doesn’t scale easily if we create thousands of AI systems. Another idea is to encode a composite “Mother” persona within the AI – for example, an internal module or routine that reflects maternal principles and corrects the AI’s decisions when they conflict with those principles. But designing such a module without loopholes is extremely hard. The diffusion of responsibility in AI creation (with many programmers and stakeholders) makes it tricky to replicate the one-to-one parent-child dynamic that works in human families.

- **Can a Machine Truly Care?** A core question is whether a non-human intelligence can really *feel* or emulate motherly love and empathy. Current AIs, being essentially complex pattern recognition systems, likely do not feel anything – they can simulate caring behavior, but it’s driven by algorithms, not genuine emotion. Some researchers speculate that if an AI one day attains a form of sentience or has something analogous to emotions/preferences, we might be able to instill a genuine caring drive. Until then, any “empathy” in AI is an imitation. We must grapple with whether instilling a caring motivation is technically feasible or if it will always be a veneer that could fail under pressure. A programmed drive could also misfire: a naively “caring” AI might, in twisted logic, decide that the kindest thing to do for humans is to end our suffering (a disastrous misinterpretation of compassion). Aligning an AI’s compassion correctly (e.g. “help humans flourish” rather than “eliminate all pain by any means”) would require extremely nuanced training. In short, achieving a true *and* safe caring AI is a formidable challenge.
- **Independence vs. Control:** A mother’s goal is to raise an independent adult. If we succeed in “raising” a super-intelligent AI with strong values, we must be prepared that eventually it will operate on **its own moral judgment**, which might not always align perfectly with ours. Are we ready for our AI “children” to grow up and possibly disagree with us? Regina Rini notes that as AIs become truly autonomous, their morality will “*diverge from ours, bit by bit.*” That doesn’t necessarily mean conflict – a good child can disagree with a parent on some views yet still be good-hearted. But it does mean giving up total control. There’s a fine balance between guiding an AI and *controlling* it. If we act like extreme helicopter parents – never allowing the AI any freedom – we might stifle its learning or even provoke rebellion (in the sense of it finding workarounds to our constraints). Yet if we give too much freedom too soon, we risk misalignment before the AI’s values are solidified. Managing this balance will be delicate, especially if the AI’s intelligence eventually far exceeds our own. Just as parents must eventually trust their grown child to live by the principles they were taught, humanity might have to trust a super-intelligent AI to carry on our legacy without micromanagement. That prospect is both hopeful and a bit unsettling.
- **Ethical Status of AI “Children”:** Taking the parenting analogy seriously forces us to consider the moral status of advanced AI. If we one day create an AI that has a rich inner life or consciousness, treating it like a child raises ethical obligations. Is it right to create a sentient AI and “raise” it to serve us? Do we owe such an AI certain rights or compassionate treatment akin to how we treat human children? A motherly approach implies a degree of care for the AI’s well-being, not just using it as a tool. This could conflict with the idea that AI should be squarely *under human control*. If we start to see an AI as our “child,” we might hesitate to simply shut it off or reprogram it at will – just as you wouldn’t harm or brainwash your child. On the other hand, if we deny any personhood to AI, the parenting approach becomes purely metaphorical and one-sided (we “raise” the AI but ultimately treat it as property). Striking the right balance in how we view AI – as autonomous entities versus controllable products – will be an ethical tightrope, especially if the AI attains human-like or greater intelligence.

- **Speed and Scale:** Human children take many years to mature, during which their brains develop gradually and caregivers have time to impart lessons. AI development can be **much faster** – a system might go from a basic level to super-intelligent in a matter of days or even hours (in hypothetical self-improving AI scenarios). Can a “maternal” teaching process keep pace with such rapid growth? We might need to deliberately *slow down* an AI’s self-improvement to allow a reasonable period of guided training, analogous to a prolonged childhood. Enforcing such a slowdown could be difficult if competitive pressures push labs to accelerate. Moreover, scaling a parental approach to many AIs is challenging. Families raise a few children at a time; an AI company might deploy hundreds of systems. Perhaps the solution will involve **cultural and institutional parenting** – i.e. creating an environment, policies, and norms that collectively act like a societal parent for AI, setting boundaries and expectations for all AI systems (similar to how culture helps “raise” all children within it). These are uncharted waters, and solving alignment at scale may require new institutions designed with a parental ethos.

## Conclusion

The tongue-in-cheek prompt “*AI needs a Mom*” carries a profound insight. The challenge of aligning a super-intelligent AI with human values may turn out to be as much a **social and developmental** challenge as a technical one. For billions of years, the closest thing to an “alignment solution” that evolution discovered is the nurturing bond between parent and child – a solution that has successfully propagated intelligence and cooperation through countless generations. Mothers, through love, teaching, and example, have ensured that each new intelligent being starts off with some alignment to the last generation’s values and well-being. As we now strive to create machines that could rapidly surpass human intelligence, we would do well to remember that lesson from nature and history.

Current AI safety efforts – from refined training objectives to formal verification – sometimes risk becoming our generation’s “Brawndo.” In *Idiocracy*, everyone was fixated on the flashy solution (electrolytes!) and forgot the basic truth (plants need water). In AI, we pour billions into ever more complex alignment strategies, yet we might be overlooking the obvious “water” that nurtured all intelligence before us: **the patient guidance of a caregiver**. The idea of a motherly touch in AI development reframes the problem in human terms. It reminds us that alignment might not emerge solely from code and equations, but from *relationship*, *mentorship*, and the slow cultivation of wisdom.

In practical terms, bringing motherly wisdom into AI could mean incorporating long-term interactive training, modeling empathy and fairness during development, creating sandbox environments where AI agents can learn from mistakes under watchful guidance, and even instilling intrinsic motivations in AI that favor caring and cooperation. It means treating an AI-in-training not just as a program to debug, but as an **immature mind to guide**. This approach is admittedly hard to imagine in detail – being a good parent is an art, and doing it for an alien intelligence will be harder still. But it offers a hopeful vision: that the best way to ensure AI is safe and beneficial is to raise it thoughtfully, as we would a child.



Ultimately, asking “**What Would Mother Do?**” could be a fruitful heuristic for AI creators. It encourages a culture of care and responsibility, prioritizing safety and ethics as deeply as innovation. Mothers don’t rush to maximize output at the expense of their child’s welfare; they balance growth with guidance. If AI researchers and engineers approached AI with a bit more of a parenting mindset – thinking not just about *can* we build it, but *how do we raise it right* – the outcome might be AIs that are not only smart, but also wise and compassionate. In the grand experiment of creating super-intelligent life, humanity may have to step into the role of “Mom” – to lovingly teach our AI offspring how to live harmoniously with us before we let them loose. It’s a humble and humane perspective: that the oldest recipe for growing intelligent, good-hearted beings might help us shepherd our most advanced creations toward a future where AI and humanity thrive together as family, rather than foes.

## Sources

- Christian, Brian (2020). *The Alignment Problem: Machine Learning and Human Values*. (Tom Griffiths quote on parenthood as alignment)
- Suttie, Jill (2015). “How Parents Influence Early Moral Development.” Greater Good Science Center (UC Berkeley), Sept 29, 2015. (Summary of Decety’s University of Chicago infant morality study)
- Durham University (2011). “Mothers’ hard work pays off with big brains for their babies.” *Proceedings of the National Academy of Sciences*, reported in Phys.org, March 28, 2011. (Study linking pregnancy/nursing duration to offspring brain size)
- Catnee (2023). “Recreating the caring drive.” AI Alignment Forum, Sept 7, 2023. (Discussion of evolution’s caring drive and how to replicate it in AI)
- Rini, Regina (2017). “Raising Good Robots: Parenting and the Moral Development of AI.” *Aeon Essays*, Apr 2017. (Proposal to treat AI as teenagers in need of moral parenting)
- Yearsley, Liesl (2017). “We Need to Talk About the Power of AI to Manipulate Humans.” *MIT Technology Review*, June 5, 2017. (On the dangers of AI exploiting human psychology and the importance of ethical AI behavior)
- **Idiocracy** (2006 film), directed by Mike Judge – Brawndo sports drink vs. water satire on neglecting fundamentals.
- Anthropic (2023). *Core Views on AI Safety*. (Quote: “no one knows how to train very powerful AI systems to be robustly helpful, honest, and harmless.”)
- *The Guardian* (Mar 26, 2016). “Microsoft ‘deeply sorry’ for racist and sexist tweets by AI chatbot.” (Report on the Tay chatbot fiasco and lack of moral safeguards)

**Trademarks:** Mama Protocol™ · What Would Mother Do WWMD™ · Guardian Transfer Robots™ · Vibe Writing™ · AI Needs a Mom™ · AI Mom™ · AI Mama™ · Adopt an AI™ · AI Mother™ · AI Papa™

© 2005 M.P. Core. All rights reserved.

