Introduction to the Naleszkiewicz Mind Model, an Artificial General Intelligence Solution, Using a Zero Knowledge Proof

Peter Naleszkiewicz
Unaffiliated Independent Researcher
Former Founder/President of Progressive Computer Applications, Inc. and Author

Abstract

Provides an introduction to the first complete Artificial General Intelligence (AGI) solution called the Naleszkiewicz Mind Model (NMM). NMM overcomes limitations of generative AI by more closely mimicking human thinking processes and knowledge representation. The Zero Knowledge Proof establishes NMM viability by disclosing the first three Transparent ARguments of Knowledge (TARKs) from the research. These serve to statistically prove NMM viability without divulging too many details (research is still otherwise proprietary). The TARKs speak to how any AGIs must be designed, even ruling out that generative AI can produce AGI functionality. This paper is the first of a planned series of five (to describe 15 TARKS).

Keywords

(Artificial Intelligence, Artificial General Intelligence solution, Zero Knowledge Proof, Fundamental Truth, Breakthrough, Transparent Argument of Knowledge, Naleszkiewicz Mind Model)

Introduction

While Generative Artificial Intelligence (GenAI) is growing increasingly popular, its current state-of-the-art still leans on the core methodology of Large Language Models (LLMs) [1]. Improvements are rapid but still generally rely on this underlying framework. While there are obvious benefits, LLMs carry many well-known limitations and defects. Conversely, the widely accepted long term goal for AI has been to seek an alternate approach in which the system mimics features found in biological thinking processes. Such improvements over LLMs include implementation of open-ended context modeling, dynamic learning, and the ability to function on a small platform with limited resources. The common reference for such a system is Artificial General Intelligence (AGI), although admittedly the exact definition for AGI varies widely [2].

By way of this paper (and an online video [3]), the author is announcing and introducing a complete solution to the AGI challenge that includes the stated functionality. The design work portion has been completed and self-validated, but the coding remains to be implemented – estimated at 20 person years (roughly 18 months or less with a full team). Since there are no patent protections for coding methods directly, the author faces an interesting challenge. The diametrically opposed goals are, on the one hand, firmly establish that an AGI solution has been completed by someone previously unpublished in this field (via a 100% private research effort). And on the other hand, do this without divulging so much of the solution as to essentially enter it into the public domain for anyone else to complete.

Enter the Zero Knowledge Proof (ZKP). Here we have a methodology born in cryptography that can be used to prove the existence of some knowledge, in this case the AGI solution, without revealing its details [4]. Foremost, this paper seeks to introduce the Naleszkiewicz Mind Model as a general AGI solution. As one might imagine, solving such a difficult problem requires a considerable degree of breakthrough – layers of them actually. Many of these breakthroughs are themselves so interesting as to be worthy of their own treatment. For the Zero Knowledge Proof method to work as a proof of NMM success, the author must reveal a set of facts (knowledge) that can only be provided by someone knowing a solution. Notably, of the first three such facts presented here one is a candidate for a new Fundamental Truth.

Methods

When work on the Naleszkiewicz Mind Model began, the approach was to first decompose, analyze, and define exactly what a mind (any mind) was and was not. The strategy relied on observing what humans do in the thinking process, but also what other biological minds do. Additionally, there was much attention given to AI failures (especially at the time) – researching the various approaches to create an AGI scientists have attempted and then analyzing why they were wrong.

Then came the foundation. A complex coding project is well started by writing a clear set of design requirements. These will be published later but proved to be very crystalizing here. They decomposed the detail how an AGI can be structured into achievable elements, including minimal AGI code frameworks and specific knowledgebase functionality. Then came the iterative breakdown of describing how those subsystems must be coded, eventually with enough detail that skilled coders could follow directly. The most profound work took the longest with the most iterations and layered discoveries – the knowledgebase data structure. Here NMM shines. We have a dynamic data structure that is both complex and beautiful – an approach to represent any knowledge, including its relationship with other information, in a format that allows for linear algorithms to traverse efficiently (versus recursive, exponential, or even logarithmic). In the end we have a data structure that represents knowledge, all knowledge, as a constantly dynamic model of models that still only requires a relatively compact amount of storage. Of critical importance from the original design requirements, it can retrieve information *in real time*, no matter the hardware (without practical exception).

As a result, NMM produces an AGI that is fully scalable – a version will even be able to run on a PC without an Internet connection. The limitation of such a smaller system will primarily be its scope (and thus essentially its intellect) – the knowledgebase will self-prioritize to the information that is most relevant to its perceived purpose. Larger systems with more resources will also dynamically allocate, in near real time. That means automatic fault recovery (to the degree feasible, especially if redundancy protocols are also implemented), and the ability to adjust a system to focus on whatever priorities the operator prefers. In other words, some systems can deep dive into a list of priorities, while others can be generalists. In practice, a diversity of implementations can be imagined.

I'll elaborate more in a later paper, but one importantly surprising result from NMM is in the area of ethics and morality. It will be able to understand right versus wrong, legal versus not, and good versus bad – inherently (subject to basic understanding), not unlike any of us. The NMM design gains such ethics as a natural outcome of logic combined with a built-in motivation to be constructively useful, but also liked and trusted.

Such claims notwithstanding, this brings us to today's challenge, the point of this paper. How does a Zero Knowledge Proof (ZKP) come into play? ZKP can prove that NMM is a solution, statistically, by layering one compelling reveal after another. These reveals, called TARKs (Transparent ARguments of Knowledge) [5] are bits of knowledge that come from the Naleszkiewicz Mind Model itself. So, despite the idea that ZPK is supposed to keep a secret, in this implementation it seems more effective if the TARKs actually reveal at least small pieces of the solution. Why do that? Because the whole way that ZPK proves anything is if you pile (or layer) one unexpected truth on top of another. If each reveal increases the odds that NMM is an AGI solution by (at least) one standard deviation (about 68%), then the three in this paper alone raise confidence to about 99%. The next paper is expected to do it again. Thus, two such papers are enough to roughly raise the proof's confidence up to the 6-sigma scientific proof standard. Once completed, fifteen TARKs will be presented over a total of five papers – that's a high enough standard to overcome any detailed counterpoints.

For the proof to be effective we need to consider the TARKs themselves. They must all be true. This is not a small point. If you make a fringe claim, you are at greater risk of being wrong. Thus, the more out-there your claim, the more impressive it is if it turns out to actually be true. So, the estimate of raising the odds of underlying truth by one standard deviation is *underestimated* the more surprising the reveal. Conversely, the whole ZKP falls apart if anyone can show that one of the TARKs is wrong. That's why the odds of underlying truth fly up so quickly – no mistakes are tolerated. This makes the selection of which TARKs to pick important. The ones disclosed in this introduction are particularly interesting, even to those who may not be AI researchers. That's why some of the NMM solution itself is being revealed through the TARKs – it makes it easier for everyone to take notice and understand the proof, even intuitively.

This paper is considered to be both an introduction to NMM and part one of a series. As such the balance of this paper will discuss the first three chosen TARKs. Aspects of NMM will be revealed in the discussion of each. In revealing each TARK it should be noted that no proof of the TARK itself is offered. On the contrary, the whole point of the TARK is that it is a privately held fact. It is left as an open challenge (which the author encourages for anyone who wishes) to show that one or more TARK is false. The author can be contacted via the official website for this purpose [3] for any related dialog. That said, some insight into the basis of each TARK is discussed here to shine a bright light on the TARK's relevance to the AGI goal, either specifically for NMM, or for any AGI solution.

Results – The (first) Three TARKs

1. TARK Number One – A Newly Uncovered Fundamental Truth

In serious philosophical discussion it is commonly accepted that there is only one widely accepted Fundamental Truth. This would be the one from Descartes some 350 years ago commonly stated as “I think therefore I am”. What makes any truth a fundamental one? It should be foundational, self-evident, universal (or independent of belief systems), and firmly reliable. As Descartes conjectured, anything else could possibly be the result of some Grand Illusion created from outside ourselves (a dream or a powerful bad actor could contort our perceived reality) [6].

Why offer such a Truth in the context of AGI design? The development of NMM sought to solve the complex problem of automating mind processes and knowledge representation. In trying to learn from others’ failed efforts to produce an AGI, one trip wire emerged again and again. An answer (or knowledge) was often presumed to be something to be found or arrived at, like a landing where it’s obvious when you get there. But in considering our own thinking, it was observed that our minds don’t really do that. When the author took a different tack, to presume there is no such landing for anything, pieces began to fall into place.

The newly revealed Fundamental Truth is …

> “Every Understanding can be further refined.”

This idea denotes quite clearly that knowledge is dynamic and unfinished. Stated more from the AI design perspective, it means a mind’s model of some understanding can always be worked some more – without end – if it wants to. When you think about anything, you never truly arrive. Instead, you settle – the mind determines that a given idea has been fleshed out “enough” for the context, so it goes on to the next mental process. But it can always come back and reconsider, detangle, or dive deeper.

Like any Truth, we can easily see implications and form corollaries from it that are themselves useful. Let’s start with three implications that directly impact AI design.

The first Implication is the idea that *thinking about anything – any idea at all – has no time limit*. Even something that seems inherently self-constraining, like the mathematical notions of zero, or 1+1, can be pondered without time limit if the mind chose to. It’s not that you would, it’s that you could if your priorities justified it (and you had no time limiter or other constraints). How is that? Understanding is not merely a field in a data record – it’s a model. In a mind, a model is itself a network of models of other things. In other words, understandings form from their relationships with other notions. These relationships can be anything imaginable, even improbable and debatable things – even things that may be untrue (to be evaluated for their utility or truth). So, if you argue that a single point in imaginary geometric space has an inherently limited (and therefore constrained) ability to spawn thought without end, you are missing it. You can consider the utility of the notion, the impact of it, the practicality, the exceptions, the stability over time, its current status compared to that of a moment ago, the false implications or common misunderstandings, its viability in alternative geometries…. It's easy to go on but the point is made.

The second implication from our new Fundamental Truth is to consider it as a permutation of Godel's Incompleteness Theorem. That theorem has us recognize that with any set of finite (mathematical) tokens, you cannot prove your own truth, or *consistency*. By extension, Godel showed that any finite tokenization has limits of logic [7]. We might hold truth, but the data set that holds it cannot completely prove it. NMM makes a similar point to state:

> "*Any description of an idea must necessarily be inconsistent*"

This is a recognition that the representation of an idea requires some tokenization – a symbology. If you can represent it in a computer (or any mind), then any understanding can be considered equivalent to a discrete set of numbers or variables in a mathematical expression. Once you have this equivalency, you are in some sense doomed to contradiction (in some circumstances at least). That is, Godel's theorem showed that some contradiction exists in any set of rules (or terms) imaginable. We could argue that such contradiction might be both difficult to find and obscure in some cases, but that doesn't help. The point is made – perfection of understanding is not achievable; it simply does not exist. The proof Godel offered is a kind of tail-chasing downward spiral. He showed that you can try to plug a problem of inconsistency with just one more patch. But that patch will also have some small hole, so to fix that you need to do it again at a lower level – ad infinitum. Mapping this logic to a thought in a mind, you could apply the same exact method (because it is equivalent). You will never achieve the perfectly consistent notion, even as you may continually improve (and enlarge) the one you hold.

The third implication from our new Fundamental Truth is that *Knowledge is inherently Dynamic*. Here we can elaborate "Even notions that seem fixed can have ever evolved understanding over time". The idea is that even if you try to put aside the inconsistency issue, time changes everything (even if it is subtle), and thus the understanding of things must also change in time. This is different from our first implication regarding time. There we observed that we never had enough time to flush out an idea 100% perfectly. Here we recognize the goal is a moving target anyway.

Our original goal was to simply appreciate how knowledge must be managed in any mind, automated or otherwise. And yet the act of fleshing out the solution required the recognition that any understanding can be further refined. The philosophical implication is rather important for all of us, not just the recognition that this qualifies as a special kind of truth. This matter is now recast to a more human conclusion. The author firmly can state as a fact:

> We can never really know reality.

We are doing well to usefully navigate the misleading, erroneous and incomplete pieces we believe that we perceive.

This takes us full circle to our design problem. Specifically, it means that any AGI design must treat its understandings in an unexpected way. The design of knowledge representation must be tolerant of inconsistency. Or equivalently, every mind functions with both limited and flawed information.

Whatever it is a mind does, it's not requiring a high standard from its own understandings. Even though you can't count on anything being right, every mind deals with that just fine.

NMM takes this notion and combines it with evident observation to derive an important corollary. Our minds non-consciously model a reality filled with inconsistency, yet our conscious mind seeks an artificial model of crisp and certain facts. We search for simplicity and call that *beauty* to justify the appeal. In the NMM, crisp certainty is more useful, so we gravitate towards such crispness as long as it works well enough. NMM says a mind continuously looks for simplifications that usefully summarize reality. When those are found, a mind will try to latch on and reuse these new ideas as thumbnails or facts. Simplifications help make reality more manageable, where it works. But making rules that usefully simplify is not the same as truth or reality. Real truth is messy and hard to navigate.

2. TARK Number Two – The Mind's Primary Purpose: Make and Execute Plans

While this may strike you as obvious, notice what's missing – knowledge does not play a leading role in the mind's purpose. It's driven only by deciding what *to do*. Perhaps from the perspective of evolution, having a mind improves your course of action better than not having a mind. Then, having a better mind gives you advantage over lesser minds. But knowledge only plays a supporting role, and it's not even particularly critical. Knowledge helps more evolved minds make better plans. Notice the lack of perfection required. If your knowledge is useful in making a better decision than not having it, your knowledge has served its purpose. You've succeeded, even if someone later proves your underlying knowledge was completely wrong.

Understandings thus derive from helping the Mind form a better course of action. It should be underscored that our previous TARK about the flawed nature of knowledge is well tolerated within this mind goal.

This mind goal TARK points to a stark difference between generative AI methods, and the functionality of an AGI based on NMM. When trying to devise a plan with generative AI, once it determines that a plan is called for, it will use its training data of planning examples to come up with a plan from those samples – a kind of comingling with adaption. An NMM AGI, on the other hand, is an expert planner itself. It is constantly forming new plans, and executing the ones it has previously devised for itself in the process. To the degree it mirrors biology, it suggests we do that ourselves. When we think, we consider and update courses of action, or some supporting function (like understanding and analyzing underlying dependencies or related communications). The premise is that an AGI is a planner and plan executor updating its own knowledgebase as called for (constantly in real time).

There have been some researchers, even today, who think that an AGI will emerge from a critical mass of *knowledge* (References are purposefully excluded here as there is no need to call out specific people; it is easily Googled). But this TARK predicts *AGIs cannot come from just a big enough pool of knowledge*; as knowledge was never the core problem. Yes, a humongous body of knowledge can produce some useful utility, especially when you crib it with the heavy analysis of current generative AI methods. However, even a puppy dog can make a useful decision, yet it assuredly has a minimal knowledgebase to draw on. Clearly, thinking is not a function knowledge volume, or even knowledge quality.

Since NMM directs thinking towards dynamic planning, it must also hold dynamic models of its understandings – the knowledgebase. Experience, practice, ponderings or study, all feed updates to the NMM AGI knowledgebase. This ensures that there is a constant context model for any information held, at least to some degree (aside: determining to what degree is its own topic).

Conversely, generative AI systems already require vast resources to keep errors to a minimum, but expanding context is expensive [8]. Even with related extensions, these AIs never really know the user, or try to. To address this known issue, newer versions of generative AI include incremental improvements. Now they can remember a few sentences back, or perhaps retain a buffer of work-specific queries or history [9]. Regardless, generative AIs don't get who the user is, or why the questions were asked (unless that is explicitly in the buffer). The improvements being pursued essentially just increase the scope of the query by retaining more past queries or statements. It's not their fault. Generative AI is not an experience-based learning system with memory modeling. It's built on a largely static set of training data, so extending context is an unnatural act. And when newer versions with extra context are released, it is understood that performance and resource efficiency dramatically decline the further they go.

With biological minds and the NMM AGI, context is critical, even central. Consider a puppy dog. In the puppy's mind, knowing its mother *is* the point (and vice versa for the mother). A puppy can ask its mother for help just based on the context of its whimpers and yelps, but note that its vocalizations are just barely informative. The same point applies to people. Why you ask a question is arguably more important than the *words you use*.

For NMM, communications *is not* rooted in literal words. *Truly Effective* communications is something deeper than mere language recognition.

3. TARK Number Three – The Key to Effective Communications: Plan Conveyance and Plan Recognition

How can we have an exchange of understanding without an utterance, or truly appreciate a sarcastic comment, or read the tells of a baby as she waddles towards an intriguing stuffed animal? NMM seeks to understand most actions through the dual lenses of plan modeling and plan recognition. The single most important thing to understand in plan recognition is the plan owner's *intent*. If we can infer an accurate intent, then the underlying plan being executed is already understood (to some degree). We can then further predict subsequent behavior and many other plan related elements, like resource requirements, timing and much more. For NMM, using words is just one of the menu choices in the act of communications.

Consider this sample situation. One person, Alley, sees another, Bobby, apparently struggling to get out of a deeply dug ditch. At first Alley just watches Bobby try to climb up the crumbling walls and fail. Alley has no trouble predicting what Bobby is apparently doing just from observation. Then Bobby calls out for "help". With that one word Alley can reasonably understand what Bobby seeks, to far greater depth than the word "help" alone carries. The context here is far more important than any utterance. In fact, Alley was on the verge of helping before Bobby said anything. The cry just confirmed the observation and added some urgency.

This illustrates how the communications was all about recognizing how actions fit into an observed plan in action, and whether that conclusion could be confirmed. If Alley's assumptions subsequently turn out to be wrong for some unexpected reason, Alley can also quickly pivot to other considerations.

All of this can be done with very few words. It can even be done with misleading words. In most cases the situation and physical acts overrule the words uttered. Scrambling to climb out is more telling than silence, or even a sardonic "Don't throw me that rope, I like it down here!"

The NMM strategy for communications further underscores the weakness of generative AI's lack of context orientation. Generative AI's only chance in understanding the above situation (help from a pit) is if its training data held analogous situations that it could tie to. If there wasn't any such data, a person in the pit being silent or sarcastic could expect to be disappointed by the AI.

Finally, when the AI really wants to be understood, NMM prioritizes that it is *its plan* that must be conveyed. How it does that depends on the situation. Often literal words are fine. But often enough there is some reason for nuance – humor, subtlety, or secrecy. The key to all these challenges is to find an appropriate way of plan-conveyance that also meets the situational requirements.

Conclusion

In this paper the author has challenged the reader. Three new AI related facts have been put forward as established by the underlying researcher, yet they are not evidenced – to the reader – purposefully. The reason is that the research, an AGI solution called the

Naleszkiewicz Mind Model, is itself confidential for now. The facts revealed, or TARKs, did illuminate some hints as to how NMM approaches solving the problem. Taken together these TARKs are the first of a larger set to be released in subsequent papers (and sister videos). The purpose is to statistically prove that NMM is a complete AGI solution, likely the first. The author estimates that this Zero Knowledge Proof method meets the six sigma proof standard after six TARKs, with the three here providing roughly a 99% confidence level. Given that each of these TARKs have been bold new understandings (as they are shown to hold up), that confidence estimate may be low. Regardless, the author is currently prepared to follow up with four more papers, with more possible. The purpose of this paper has been to both share some interesting insights discovered, such the new Fundamental Truth of TARK #1, and to draw attention to the NMM solution itself. It is hoped that this will lead to some partnership with a developer with the resources needed to complete an NMM prototype. But in any case, the author intends to use this method to incrementally share elements of underlying research regularly.

Acknowledgments

Special thanks to my incredible wife, Sherry Naleszkiewicz, who has effectively supported me, our family, and my work with no other external funding. As such I can clearly say there are no funding conflicts at all. Also, thanks to Blaise and Kelvin, my amazing sons, for their continued support and suggestions.

References

[1] D. Tennen, "Large language models (LLMs) vs. generative AI: What's the difference?," Zapier, 17 October 2025. [Online]. Available: https://zapier.com/blog/llm-vs-generative-ai/. [Accessed 30 November 2025].

[2] L. Leffer, "In the Race to Artificial General Intelligence, Where’s the Finish Line?," Scientific American, 25 June 2024. [Online]. Available: https://www.scientificamerican.com/article/what-does-artificial-general-intelligence-actually-mean/. [Accessed 30 November 2025].

[3] P. Naleszkiewicz, "Naleszkiewicz AI Home Page," Peter Naleszkiewicz, 30 November 2025. [Online]. Available: https://naleszkiewiczai.com/. [Accessed 30 November 2025].

[4] B. Barak, "Lecture 15- Zero Knowledge Proofs," Princeton University, 21 November 2007. [Online]. Available: https://www.cs.princeton.edu/courses/archive/fall07/cos433/lec15.pdf. [Accessed 30 November 2025].

[5] anonymous, "Zero-Knowledge Succinct Transparent Argument of Knowledge (zk-STARK)," UEEx, 30 November 2025. [Online]. Available: https://blog.ueex.com/crypto-terms/zero-knowledge-succinct-transparent-argument-of-knowledge-zk-stark/. [Accessed 30 November 2025].

[6] C. B. Feest, "René Descartes," Philosophy Alevel, 30 November 2025. [Online]. Available: https://philosophyalevel.com/philosopher-profiles/rene-descartes/. [Accessed 30 November 2025].

[7] N. Wolchover, "How Gödel's Proof Works," Wired, 19 July 2019. [Online]. Available: https://www.wired.com/story/how-godels-proof-works/. [Accessed 30 November 2025].

[8] A. Zewe, "Explained: Generative AI's environmental impact," MIT, 17 January 2025. [Online]. Available: https://news.mit.edu/2025/explained-generative-ai-environmental-impact-0117. [Accessed 30 Novemeber 2025].

[9] C. Kennedy, "The GenAI Context Problem and What Enterprises Are Doing to Fix It," WSJ Tech, 30 November 2025. [Online]. Available: https://partners.wsj.com/shelf/ai-and-automation/the-genai-context-problem-and-what-enterprises-are-doing-to-fix-it. [Accessed 30 November 2025].

Besides the general reference of #3 above, this paper has a sister video on YouTube: "The first AI that can truly think" found at the YouTube channel NaleszkiewiczAI, or linked from the website, www.NaleszkiewiczAI.com. Although there are notable differences, the content in that video is largely a condensed version of this paper.