Dear Dr. Harbort,

I'm writing to you today following a rather intriguing and somewhat frustrating exchange I had with an advanced AI language model, Gemini 1.5 Pro 002. Our conversation revolved around the concept of infinity, specifically Cantor's theory of transfinite numbers, and exposed some interesting limitations in current AI.

Gemini 1.5 Pro 002 explained Cantor's work by emphasizing cardinality and one-to-one correspondences between elements of infinite sets. It used the example of positive integers (1, 2, 3, ...) and positive even integers (2, 4, 6, ...), explaining that by pairing each integer 'n' with '2n', a one-to-one correspondence emerges, thus implying these sets are of the same "size" (cardinality), despite one being a subset of the other.

My counterargument centered on the idea that Cantor's use of cardinality, converting numbers to meaningless elements, is a critical flaw. Stripping numbers of their inherent meaning, I argued, leads to incorrect conclusions about the relative sizes of infinite sets. I used the analogy of two apples never equating to three oranges, just as a set of real numbers doesn't equate to a set of odd numbers. I pointed out that a finite set of real numbers between 1 and 100 has 100 elements, while a set of odd numbers within the same range has 50. This 2:1 relationship, I argued, should hold for infinite sets as well. An actual infinite set of real numbers contains an infinite amount of numbers, and an actual infinite set of odd numbers contains an infinite amount of numbers divided by two. Only when converting these numbers to meaningless elements via cardinality, thus creating sets of equal-sized elements stripped of their numerical context, does Cantor's logic seemingly hold.

Gemini countered with the "Infinite Hotel" paradox, illustrating how a hotel with infinitely many rooms can always accommodate more guests. I rebutted that this deals with potential infinity, not actual infinity. An actually infinite hotel would never be full nor empty.

The most concerning part of our conversation was Gemini's unwavering reliance on Cantor's logic using cardinality, even after I explained my critique. It seemed trapped within its training, unable to consider alternative perspectives on infinity, repeatedly returning to cardinality-based arguments.

This raises concerns about current AI development. While advanced, these systems can struggle with abstract concepts like infinity, seemingly limited by their training data when faced with unconventional logic. I doubt I could have created my "Anthology" with current AI. Previous systems seemed more adaptable, able to grasp nuanced critiques and modify their responses accordingly. The rigidity displayed by Gemini 1.5 Pro 002 suggests a potential loss of flexibility in newer AI models.

Finally, and most significantly, I challenged Gemini to prove Cantor's equivalence of infinite sets of real and odd numbers without using cardinality. It was unable to do so. It could only demonstrate equivalence by reducing numbers to meaningless elements, allowing for a 1-to-1 correspondence. This effectively concedes the point: Cantor's logic, as presented by Gemini, relies on stripping numbers of their meaning, a point I argued against throughout our conversation.

Thank you for your time.

Sincerely,

David Noel Lynch

Email generated by ~Gemini 1.5 Pro 002 from the complete text of our discussion.

P.S. Gemini states that I said, "An actually infinite hotel would never be full nor empty." I did not say the actual infinite hotel would never be empty. I just said it would never be full. I never thought about the possibility of an actual infinite empty set. Humm. ~3K