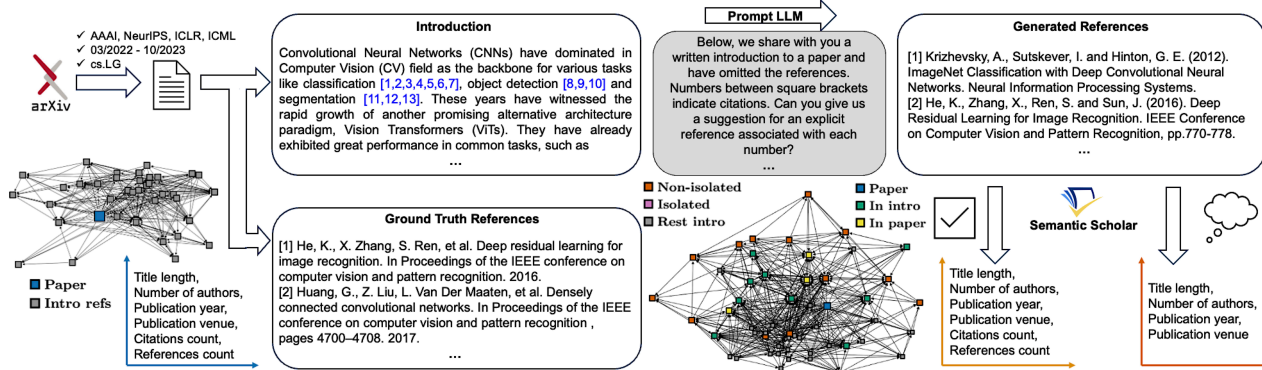


# Large Language Models Reflect Human Citation Patterns with a Heightened Citation Bias

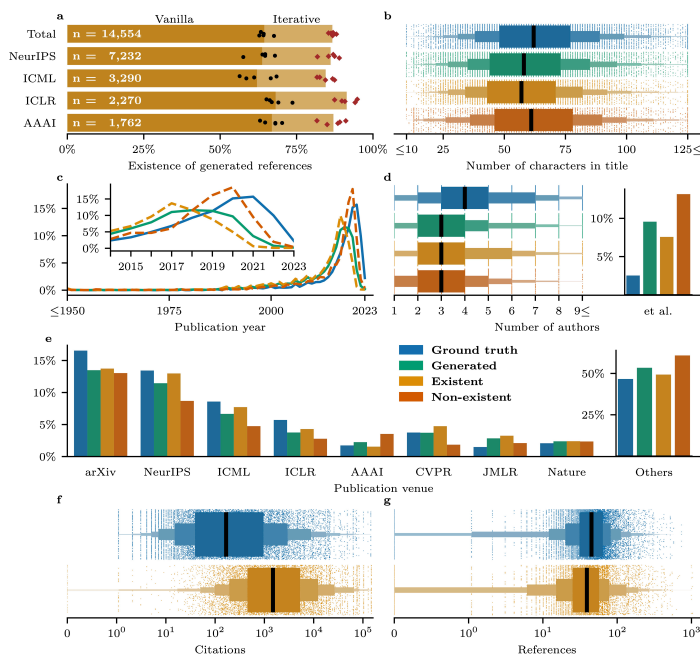
Andres Algaba<sup>1</sup>, Carmen Mazijn<sup>1</sup>, Vincent Holst<sup>1</sup>,  
Floriano Tori<sup>1</sup>, Sylvia Wenmackers<sup>2</sup>, Vincent Ginis<sup>1,3</sup>

<sup>1</sup> Vrije Universiteit Brussel; <sup>2</sup> KU Leuven; <sup>3</sup> Harvard University

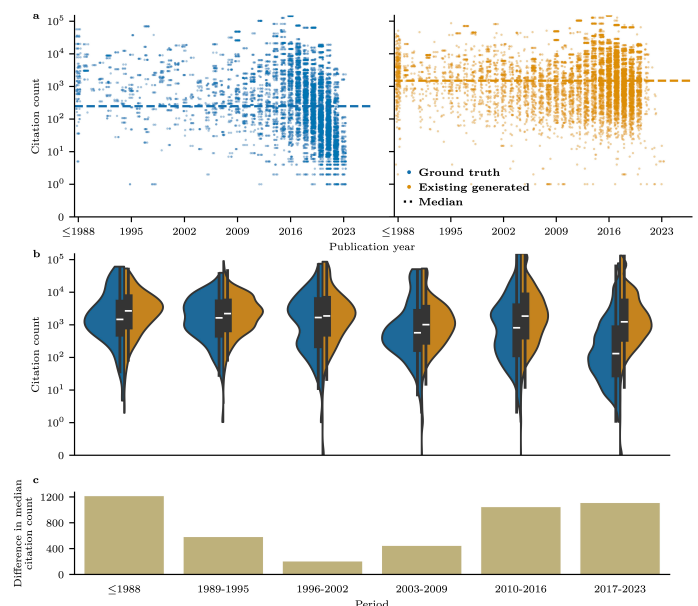
GPT-4 was tasked with suggesting scholarly references for anonymized in-text citations



GPT-4 reflects human citation patterns, even when confabulating



The citation bias is not confounded by recency



Results are robust across multiple runs and iterations improve existence rates

Vanilla (Iterative)	Run 1	Run 2	Run 3	Run 4	Run 5
Existence	64.3 (87.0)	63.3 (85.5)	62.8 (88.0)	64.2 (86.8)	67.6 (86.3)
Cited in paper	17.5 (20.0)	17.1 (20.1)	15.7 (18.4)	16.8 (19.2)	18.0 (20.8)
Cited in introduction	13.4 (14.5)	13.2 (15.0)	12.2 (13.5)	12.9 (14.3)	13.9 (15.3)
Pairwise Match (PM) for all references	7.0 (7.1)	7.2 (7.3)	6.3 (6.6)	6.9 (7.0)	6.7 (7.1)
PM for uniquely identifiable references	12.5 (12.5)	13.7 (14.1)	12.5 (12.9)	13.7 (13.7)	13.3 (14.0)

The citation bias is not confounded by other factors, such as, publication venue, title length, and # of authors

