

Peer Review

Review of: "SafeSynthDP: Leveraging Large Language Models for Privacy-Preserving Synthetic Data Generation Using Differential Privacy"

Anand Vunnam¹

1. Independent researcher

I found this paper fascinating as it tackles the challenge of balancing data privacy and utility in ML models. The integration of Differential Privacy with LLM-generated synthetic data is a smart approach, and I appreciate the thorough evaluation of its effectiveness. It's a promising step toward privacy-compliant AI research.

Declarations

Potential competing interests: No potential competing interests to declare.