

## **Making GenAI Greener & Sustainable**

The significant upscaling of large language models (LLMs) in recent years has increased demands for computing power and energy, raising concerns about environmental impact and sustainability. In this work, we propose a lightweight compression method based on Singular Value Decomposition (SVD), guided by mutual information (MI), to reduce model size and resource usage while maintaining performance. We apply our method to a well-known open source AI from Meta. Our experiments demonstrate that our method reduces memory and computation requirements without compromising performance. These results indicate that our method effectively reduces energy consumption. We also calculate the Co2 footprint of GenAI.