



# Accelerating Policy Decisions: A Multi-Agent System for Interactive Question- Answering on Climate Policy Data

Elizabeth Osanyinro<sup>1</sup>, Ernest Effiong Offiong<sup>2</sup>, Maxwell Nwanna<sup>3</sup>, Grace Farayola<sup>4</sup>  
Olaitan Olaonipekun<sup>5</sup>, Abiola Oludotun<sup>6</sup>, Sayo Agunbiade<sup>7</sup>, Oladotun Fasogbon<sup>8</sup>  
Ogheneruona Maria Esegbona-Isikeh<sup>9</sup>, Lanre Shittu<sup>10</sup>, Toyese Oloyede<sup>11</sup>, Sa'id Olanrewaju<sup>12</sup>

<sup>1</sup>University of Bradford, UK, <sup>2</sup>Carbonnote, USA, <sup>3</sup>RideNear, UK, <sup>4</sup>University of Buckingham, UK  
<sup>5</sup>Vuhosi Limited, UK, <sup>6</sup>Readrly Limited, UK, <sup>7,8,10,11</sup>Independent Researcher, UK

## ABSTRACT

Policymakers face significant challenges in extracting actionable insights from large-scale, unstructured datasets such as those generated during COP29. Traditional approaches, including fine-tuning large language models (LLMs), are computationally expensive, time-intensive, and struggle to adapt to evolving policy landscapes. A Dynamic Multi-Agent System for Policy Recommendation Optimisation is introduced, trained on 30,000 decision documents from COP29, leveraging modular agent architectures and advanced prompt engineering to generate precise, context-aware recommendations without fine-tuning. The system integrates a retrieval agent to efficiently extract relevant data from COP29 repositories using embedding-based search and vector databases, a synthesis agent that processes retrieved information through role-based, few-shot, and chain-of-thought prompting to generate structured recommendations, and a feedback agent that dynamically refines outputs based on user interactions, ensuring alignment with regional policies and evolving regulatory landscapes. This approach introduces key innovations, including prompt-driven optimisation to achieve high-quality, domain-specific outputs without retraining, multi-agent collaboration for real-time adaptability, and external knowledge integration through structured knowledge graphs and retrieval-augmented generation (RAG) to enhance factual grounding and mitigate hallucinations. Evaluation against human-expert baselines demonstrates a 34% increase in recommendation precision, a 41% reduction in processing time, and significant improvements in adaptability compared to conventional methods. A scalable, real-time, and adaptive policy recommendation process is established, advancing AI-assisted climate governance and ensuring that policy decisions remain data-driven, transparent, and actionable for global stakeholders.

**Keywords:** Climate Governance AI; Multi-Agent Systems; Policy Recommendation Optimisation; Retrieval-Augmented Generation (RAG); Sustainable Decision-Making