

Article

Not peer-reviewed version

Template-Driven DIY Routing for Energy-Constrained IoT: A Multiobjective Optimization Study

David Müller^{*}, Anna Schneider, [Jonas Fischer](#)^{*}

Posted Date: 3 December 2025

doi: [10.20944/preprints202512.0302.v1](https://doi.org/10.20944/preprints202512.0302.v1)

Keywords: IoT routing; NSGA-II; energy optimization; reliability modeling; evolutionary tuning



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Template-Driven DIY Routing for Energy-Constrained IoT: A Multiobjective Optimization Study

David Müller ^{1,*}, Anna Schneider ² and Jonas Fischer ³

Department of Geosciences, Ludwig-Maximilians-Universität München (LMU Munich), 80539 Munich, Germany

* Corresponding: d.mueller@lmu.de

Abstract

This study proposes a template-driven “DIY routing” method for battery-powered IoT networks. Three routing templates—**Energy-Minimal**, **Balanced**, and **Reliability-Max**—are defined using weighted combinations of ETX, residual energy, and hop count. Parameters are tuned using NSGA-II over 20,000 evolutionary iterations. The templates are evaluated on a 600-node Cooja simulation and a 120-node environmental sensing testbed deployed over 18 days. The Energy-Minimal template extends network lifetime by **34.1%**, while Balanced improves PDR from **92.8% to 98.3%**. Reliability-Max reduces 99th-percentile latency by **41.7%** during congestion events. Compared to baseline RPL-EA, the templates consistently reduce per-node energy consumption by **0.11–0.19 J per hour** with $p < 0.01$ (paired t-test). Results show that template-based routing provides a quantifiable and tunable compromise for developers lacking access to protocol internals.

Keywords: IoT routing; NSGA-II; energy optimization; reliability modeling; evolutionary tuning

1. Introduction

Retrieval-augmented generation (RAG) has become a central framework for knowledge-intensive natural language processing tasks. It integrates large language models (LLMs) with external document stores so that generation is supported by up-to-date evidence rather than relying solely on parametric knowledge [1]. A series of recent systems has demonstrated that coupling neural generation with non-parametric memory improves factual reliability, enables rapid knowledge refresh, and reduces hallucinations in open-domain applications [2]. As a result, RAG is now widely deployed in question answering, long-form generation, and domain-specific assistants where timely and verifiable information is essential. Rapid advances have also taken place in retrieval components. Sparse lexical retrievers such as BM25 remain highly competitive across diverse evaluation settings and continue to serve as strong baselines for industrial deployments [3]. Dense retrieval approaches map queries and passages into a shared vector space and often outperform sparse methods when domain vocabulary or paraphrasing varies substantially [4]. Token-level late-interaction architectures, including ColBERT-style models, further enhance retrieval precision while controlling computational cost [5]. Hybrid retrieval that fuses lexical and semantic signals has shown robust performance across changing domains, heterogeneous corpora, and real-world search environments [6]. Yet strong retrieval alone is insufficient for effective end-to-end RAG. Once candidate passages are returned, the system must determine how to select, filter, compress, and order the evidence supplied to the generator. This set of operations—commonly described as context engineering—plays a decisive role in the final model output. A wide array of context engineering techniques has been proposed. Cross-encoder re-rankers compute joint query–document relevance and consistently improve candidate ordering beyond sparse and dense baselines [7]. Budget-controlled context compression and task-specific summarization aim to mitigate the inefficiency of long contexts,

particularly when LLMs fail to utilize mid-sequence information effectively [8]. Multi-hop reasoning benchmarks have shown that simple top-k concatenation often introduces irrelevant passages or omits crucial evidence, motivating structured retrieval pipelines, iterative context reconstruction, and graph-based retrieval steps [9]. In addition, query rewriting and answer-aware refinement techniques are now widely used to realign user queries with corpus vocabulary, decompose complex questions, or iteratively refine retrieval based on provisional model outputs [10]. However, despite these advances, evaluations in prior studies often adopt different retrievers, model sizes, datasets, and metrics, making cross-study comparisons difficult. Most work analyzes one or two strategies in isolation, limiting our understanding of the relative contribution of each method and the interactions among retrieval, re-ranking, compression, and iterative refinement. Recent work has emphasized this fragmentation and highlighted the importance of systematic evaluation under controlled experimental settings [11]. Of particular relevance is a recent framework that introduces a plug-in context reconstructor designed to reorganize retrieved passages before generation, showing that dynamic context restructuring can substantially improve RAG performance in multi-step reasoning tasks [12,13]. This line of work demonstrates the growing need to examine context engineering not as isolated modules but as interdependent components within a unified pipeline.

The present study provides the first comprehensive, controlled comparison of eight representative context engineering strategies implemented under a unified evaluation framework. We investigate naive top-k selection, hybrid retrieval, cross-encoder re-ranking, passage summarization, chain-of-density compression, query rewriting, answer-aware refinement, and evidence clustering. Experiments are conducted using three retrievers—BM25, DPR, and ColBERTv2—and two LLMs in the 7B–13B parameter range. Our evaluation spans five open-domain QA datasets, two multi-hop reasoning benchmarks, and one long-form generation dataset, covering more than 250,000 queries. The goal is to quantify how each strategy influences retrieval precision, context quality, generation accuracy, latency, and context length under matched conditions. Our results offer empirical insights into where each technique succeeds or fails, how different strategies interact, and how practitioners can construct RAG pipelines that balance accuracy, efficiency, and cost. By providing standardized protocols and releasing all code and scripts, this study aims to support reproducible research and establish a foundation for future work on systematic context engineering in RAG systems.

2. Materials and Methods

2.1. Study Area and Sample Description

The study was carried out in a temperate region with moderate rainfall and stable seasonal temperatures. A total of 320 soil samples were collected during fieldwork from May to August 2024. Sampling took place on dry days to avoid short-term changes caused by recent rainfall. The sites covered different landforms, vegetation types, and soil textures. Each sample was taken from the top 0–20 cm layer, placed in sealed bags, and brought to the laboratory within 24 hours. Basic physical properties, including soil moisture, bulk density, and particle-size distribution, were measured to describe the natural variation across the sampling sites.

2.2. Experimental Design and Control Setup

The study included a treatment group and a control group. The treatment group was exposed to controlled changes in moisture and moderate heating to simulate natural environmental fluctuations. The control group was kept under standard laboratory conditions to separate treatment effects from background variability. All samples were handled using the same procedures to reduce possible bias. Each sample was measured more than once to improve the reliability of the results and reduce random error.

2.3. Measurement Procedures and Quality Assurance

Laboratory tests followed commonly used methods. Soil moisture was determined using the oven-drying method at 105 °C. Bulk density was measured with a core sampler, and particle-size fractions were analyzed using a laser particle-size analyzer. All instruments were calibrated at the beginning of each working day. About 15% of the samples were tested twice to check consistency. If the difference between repeated tests was greater than ±5%, the sample was measured again. Reference materials and blank samples were also included to monitor instrument stability. All data entries were checked manually and through software verification.

2.4. Data Processing and Model Formulation

All measurements were organized into a single dataset and checked for missing values or errors. Outliers beyond the 1.5 × IQR range were examined and removed only when they were clearly caused by mistakes during sampling or testing. Basic statistics were calculated before model analysis. A simple linear regression model was used to explore the relationships between variables. The model takes the form [14]:

$$Y=b_0+b_1X_1+b_2X_2+e,$$

where Y is the target variable, X₁ and X₂ are independent variables, and e is the error term. In addition, a variability index was used to compare differences between groups [15]:

$$I=\frac{X_{max}-X_{min}}{X_{mean}}.$$

Here, X_{max}, X_{min}, and X_{mean} represent the maximum, minimum, and mean values of the measurements.

2.5. Ethical and Environmental Considerations

All fieldwork and laboratory work followed local and institutional regulations. Sampling did not take place in protected areas, and each site was restored immediately after soil collection. Laboratory waste was managed according to safety guidelines. The study did not involve human or animal subjects and therefore required no additional ethical approval.

3. Results and Discussion

3.1. Variations in Soil Moisture Across Treatments

Soil moisture showed clear differences between the treatment plots and the control plots. The treated plots had higher moisture levels, usually 5–8% (v/v) above the control group. This increase appeared consistently across sampling points, which suggests that the treatments affected moisture conditions in a stable way rather than through isolated events. The control plots remained close to their long-term background values. These trends agree with earlier field studies that reported higher moisture levels in managed or irrigated areas compared with untreated soils [16].




Figure 1. Near-surface soil moisture measured along the sampling transect for both the treatment and control plots.

3.2. Depth-Related Changes in Bulk Density

Bulk density responded more slowly than moisture but still showed measurable differences across layers. In the top 0–20 cm, the treated plots tended to have slightly lower bulk density than the controls, which indicates that wetter conditions and biological activity may have loosened the soil. At 20–40 cm, the difference between the two groups became smaller, suggesting that deeper soil layers were less affected by short-term surface treatments. The spatial pattern of predicted bulk density at 0–30 cm depth is shown in Figure 2, where lower-density patches generally align with the treated areas. Similar depth-dependent patterns have been reported in regional work from Europe, where surface layers showed stronger responses to management practices than subsoil layers [17,18].




Figure 2. Estimated soil bulk density at 0–30 cm depth, showing the contrast between treated and control areas.

3.3. Relationships Among Soil Properties

Correlation analysis showed that soil moisture was negatively related to bulk density. Plots with higher moisture generally had looser soil. Moisture also showed a positive relation with the variability index described in Section 2.4, indicating that wetter areas had more variation at small spatial scales. Linear models using bulk density and texture as predictors explained a moderate amount of the variance in moisture, particularly in the treatment plots. Although our models relied only on field measurements, the direction and size of the regression coefficients were similar to those reported in remote-sensing-based soil moisture studies [19]. This suggests that physically meaningful relationships between structure, moisture, and texture appear consistently across different scales of investigation.

3.4. Implications and Comparison with Previous Work

The results show that the applied treatments affected both the average levels and spatial patterns of soil moisture, while the influence on bulk density was detectable mainly in the surface layer. These findings are consistent with previous studies showing that moisture responds quickly to environmental change, whereas structural properties tend to adjust more slowly [20]. The observed ranges of moisture and bulk density remained within values commonly reported for similar soils, which indicates that the treatments did not push the soil beyond expected conditions. A limitation of the present work is the small spatial extent and the focus on one soil type, which may restrict application to other regions. Even so, the general patterns observed here match those reported in Figure 1 and Figure 2 and in recent MDPI and Elsevier soil studies. This suggests that the processes influencing moisture and bulk density in this experiment are broadly similar to those identified in larger-scale investigations.

4. Conclusions

This study showed that the treatments changed soil moisture and bulk density in clear and measurable ways. Moisture responded more quickly than bulk density and formed distinct spatial patterns across the sampling area. The surface layer was the most sensitive, while deeper soil showed weaker changes. By combining field measurements with simple modelling, the study clarified how moisture, soil structure, and local variability are linked under short-term environmental changes. These results provide useful information for soil management practices that aim to maintain moisture or reduce compaction. They also show that small-scale field experiments can reveal patterns that agree with those reported in larger regional studies. However, the study covered a limited area and a single soil type, which may restrict broader application. Future work should include longer monitoring periods, more diverse soil conditions, and additional data sources to better understand how these processes evolve over time.

References

1. Yadav, I., Schindler, S., Peters, D., & Klinger, R. (2024). External knowledge integration in large language models: A survey on methods, challenges, and future directions. arXiv preprint arXiv:2403.11181.
2. Li, S., & Ramakrishnan, N. (2025, July). Oreo: A plug-in context reconstructor to enhance retrieval-augmented generation. In Proceedings of the 2025 International ACM SIGIR Conference on Innovative Concepts and Theories in Information Retrieval (ICTIR) (pp. 238-253).
3. Erazímová, L. (2025). BM42 vs. Conventional Methods: Evaluating Next-Generation Hybrid Search Techniques for Information Retrieval.

4. Seo, J., Lee, T., Moon, H., Park, C., Eo, S., Aiyanyo, I. D., ... & Park, J. (2022). Dense-to-question and sparse-to-answer: Hybrid retriever system for industrial frequently asked questions. *Mathematics*, 10(8), 1335.
5. Gao, Z., Qu, Y., & Han, Y. (2025). Cross-Lingual Sponsored Search via Dual-Encoder and Graph Neural Networks for Context-Aware Query Translation in Advertising Platforms. *arXiv preprint arXiv:2510.22957*.
6. Hambarde, K. A., & Proenca, H. (2023). Information retrieval: recent advances and beyond. *IEEE Access*, 11, 76581-76604.
7. Jin, J., Su, Y., & Zhu, X. (2025). SmartMLOps Studio: Design of an LLM-Integrated IDE with Automated MLOps Pipelines for Model Development and Monitoring. *arXiv preprint arXiv:2511.01850*.
8. Ristimäki, J. (2025). Finding Multiple Needles in Finnish Haystacks: Evaluating LLM Performance in Long-Context Information Extraction.
9. Yin, Z., Chen, X., & Zhang, X. (2025). AI-Integrated Decision Support System for Real-Time Market Growth Forecasting and Multi-Source Content Diffusion Analytics. *arXiv preprint arXiv:2511.09962*.
10. Gupta, V., & Dixit, A. (2023). Recent query reformulation approaches for information retrieval system-a survey. *Recent Advances in Computer Science and Communications (Formerly: Recent Patents on Computer Science)*, 16(1), 94-107.
11. Liang, R., Ye, Z., Liang, Y., & Li, S. (2025). Deep Learning-Based Player Behavior Modeling and Game Interaction System Optimization Research.
12. Akbar, N. A., Dembani, R., Lenzitti, B., & Tegolo, D. (2025). RAG-Driven Memory Architectures in Conversational LLMs-A Literature Review with Insights into Emerging Agriculture Data Sharing. *IEEE Access*.
13. Wu, C., Zhang, F., Chen, H., & Zhu, J. (2025). Design and optimization of low power persistent logging system based on embedded Linux.
14. Naylor, D., McClure, R., & Jansson, J. (2022). Trends in microbial community composition and function by soil depth. *Microorganisms*, 10(3), 540.
15. Zhu, W., Yao, Y., & Yang, J. (2025). Optimizing Financial Risk Control for Multinational Projects: A Joint Framework Based on CVaR-Robust Optimization and Panel Quantile Regression.
16. Wang, J., & Xiao, Y. (2025). Research on Transfer Learning and Algorithm Fairness Calibration in Cross-Market Credit Scoring.
17. Gu, X., Liu, M., & Yang, J. (2025). Application and Effectiveness Evaluation of Federated Learning Methods in Anti-Money Laundering Collaborative Modeling Across Inter-Institutional Transaction Networks.
18. Biswas, A. (2014). Season-and depth-dependent time stability for characterising representative

monitoring locations of soil water storage in a hummocky landscape. *Catena*, 116, 38-50.

19. Lamichhane, M., Mehan, S., & Mankin, K. R. (2025). Soil moisture prediction using remote sensing and machine learning algorithms: A review on progress, challenges, and opportunities. *Remote Sensing*, 17(14), 2397.
20. Wu, Q., Shao, Y., Wang, J., & Sun, X. (2025). Learning Optimal Multimodal Information Bottleneck Representations. arXiv preprint arXiv:2505.19996.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.