

A Demonstration of GPTUNER: A GPT-Based Manual-Reading Database Tuning System

Jiale Lao*
Sichuan University
Chengdu, China
solidlao.jiale@gmail.com

Yibo Wang*
Sichuan University
Chengdu, China
wangyibo.cs@gmail.com

Yufei Li*
Sichuan University
Chengdu, China
liyufeievangeline@gmail.com

Jianping Wang
Northwest Normal University
Lanzhou, China
2022222119@nwnu.edu.cn

Yunjia Zhang
University of Wisconsin-Madison
Madison, Wisconsin, USA
yunjia@cs.wisc.edu

Zhiyuan Cheng
Purdue University
West Lafayette, Indiana, USA
cheng443@purdue.edu

Wanghu Chen
Northwest Normal University
Lanzhou, China
chenwh@nwnu.edu.cn

Yuanchun Zhou
Computer Network Information
Center, Chinese Academy of Sciences
Beijing, China
zyc@cnic.cn

Mingjie Tang[†]
Sichuan University
Chengdu, China
tangrock@gmail.com

Jianguo Wang
Purdue University
West Lafayette, Indiana, USA
csjgwang@purdue.edu

ABSTRACT

CCS CONCEPTS

• Information systems → Database administration.

KEYWORDS

Database Tuning, Large Language Model, Bayesian Optimization

ACM Reference Format:

Jiale Lao, Yibo Wang, Yufei Li, Jianping Wang, Yunjia Zhang, Zhiyuan Cheng, Wanghu Chen, Yuanchun Zhou, Mingjie Tang, and Jianguo Wang. 2024. A Demonstration of GPTUNER: A GPT-Based Manual-Reading Database Tuning System. In *Companion of the 2024 International Conference on Management of Data (SIGMOD-Companion '24)*, June 9–15, 2024, Santiago, Chile. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3626246.3654739>

1 INTRODUCTION

Tuning configurable parameters (i.e., knobs) of modern Database Management Systems (DBMS) is crucial to improve performance. Since it is challenging to manage hundreds of knobs under diverse

database instances and query workloads in the changing cloud environment, state-of-the-art methods turn to machine learning (ML) techniques to automate the tuning process. However, they either still incur significant tuning costs because they only rely on the runtime feedback of benchmark evaluations [6, 7], which is inefficient when the search space is high-dimensional (e.g., DBMS expose hundreds of knobs) and heterogeneous (e.g., a knob can be in continuous or categorical values), or they only yield sub-optimal performance since they utilize domain knowledge in a limited way (e.g., only consider suggested values of knobs from manuals) [4, 5].

Unlike ML-based tuning methods that tune DBMS based on performance statistics, human database administrators (DBAs) leverage domain knowledge instead, including which knobs are worth tuning under certain scenarios¹, and which values to be considered given the unique semantics of a knob². While it is widely acknowledged that such knowledge is invaluable in guiding the tuning process, such wisdom seems exclusive to humans and inaccessible to machines due to the barriers in natural language understanding.

Recently, the advent of LLM makes it possible to leverage natural language knowledge. However, even equipped with the powerful LLM, it is still non-trivial to bridge the gap between black-box optimization and white-box domain knowledge, mainly due to three challenges: *C1. Lengthy and complex data pipeline*. Domain knowledge typically comes in the form of DBMS documents and discussions from web forums, which could be heterogeneous in format and noisy in contents. To leverage such knowledge, it involves a complex and lengthy data pipeline: data ingestion, data cleaning,

*The authors contributed equally to this paper.

[†]The corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGMOD-Companion '24, June 9–15, 2024, Santiago, Chile.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0421-5/24/06
<https://doi.org/10.1145/3626246.3654739>

¹For example, for an IO-intensive OLTP workload, it is recommended to adjust knob “effective_io_concurrency” from PostgreSQL.

²For instance, we should set “random_page_cost” to 1.x if we are using SSD disks.

data integration and data extraction. *C2. The brittle nature of LLM.* It is challenging to utilize LLM to solve complex and domain-specific tasks, because small modifications to the prompt can cause large variations in the model outputs, and the inevitable hallucination problem of LLM (i.e., LLM can generate answers that seem correct but are factually false). *C3. Lack of a knowledge-aware optimization framework.* The inherent design of optimization algorithms like Bayesian Optimization (BO) and Reinforcement Learning (RL) does not support the integration of external domain knowledge directly, necessitating extensive modifications to their standard workflows.

Our recent work GPTUNER [3], which is accepted by VLDB 2024, addresses these challenges using an LLM-enhanced BO approach. For *C1* and *C2*, we develop a LLM-based pipeline with two error correction mechanisms to collect and refine heterogeneous knowledge, and propose a prompt ensemble algorithm to unify a structured view of the refined knowledge. For *C3*, using the structured knowledge, we design a workload-aware and training-free knob selection strategy, develop a search space optimization technique considering the value range of each knob, and propose a Coarse-to-Fine Bayesian Optimization Framework to explore the optimized space.

This demo lets conference attendees experience GPTUNER in actions. After specifying the target DBMS, the optimization metric and the customized workload, visitors can (1) invoke an LLM to help identify which knobs are worth tuning and only tune these knobs, (2) interact with an LLM-based pipeline to see how multi-source heterogeneous knowledge with noise is transformed into a unified structured view, (3) visualize the process of using the structured knowledge to optimize the value range of each knob, (4) execute the BO to explore the optimized space, and finally (5) visualize the optimization process and export the satisfied knob configuration in SQL format such that it can be deployed directly. Our implementation and a video are available at <https://github.com/SolidLao/GPTuner> and <https://youtu.be/Hz5Zck-9TIA>, respectively.

2 SYSTEM OVERVIEW

Workflow. Figure 1 depicts the tuning pipeline of GPTUNER. ❶ User provides the DBMS to be tuned (PostgreSQL or MySQL), the target workload (TPC-C, TPC-H or a customized workload), and the optimization objective (latency or throughput). ❷ GPTUNER collects and refines the heterogeneous knowledge from different sources (GPT-4, DBMS manuals and web forums) to construct *Tuning Lake*, a collection of DBMS tuning knowledge. ❸ GPTUNER unifies the refined tuning knowledge from *Tuning Lake* into a structured view accessible to machines (e.g., JSON). ❹ GPTUNER reduces the search space dimensionality by selecting important knobs to tune (i.e., fewer knobs to tune means fewer dimensions). ❺ GPTUNER optimizes the value range of each knob based on structured knowledge. ❻ GPTUNER explores the optimized space via a novel Coarse-to-Fine Bayesian Optimization framework, and finally ❼ identifies satisfactory knob configurations within resource limits (e.g., the maximum optimization time or iterations specified by users).

Components. We briefly discuss the three components of GPTUNER in the following sections. For more technical details and experimental results, please refer to our research paper [3].

3 KNOWLEDGE HANDLER

3.1 Knowledge Preparation

In this part, *Knowledge Handler* aims to prepare a reliable natural language knowledge base. It leverages LLM to (1) prepare multi-source knowledge from GPT-4, DBMS manuals and web forums, (2) filter out noisy contents by comparing the knowledge with its DBMS system view (e.g., *pg_settings* from PostgreSQL), (3) summarize the multi-source knowledge by handling the possible conflict between the remaining knowledge in a priority way (i.e., manuals have the highest priority while GPT-4 has the lowest due to its hallucination problem), and (4) check and revise the summarization to make sure it is factual consistent with the source contents.

3.2 Knowledge Transformation

In the knowledge transformation stage, *Knowledge Handler* intends to transform the refined knowledge into a structured view accessible to machines (e.g., JSON). In the context of DBMS knob tuning, we primarily consider four types of attributes: *suggested_value*, *min_value*, *max_value* and *special_value*, whose roles are discussed in Section 4.2. Next, we use LLM to extract the values of the four attributes from the knowledge. To address the effect of the brittle nature of LLM as much as possible, we propose a *Prompt Ensemble Algorithm* since it is useful to acquire a more reliable result by aggregating multiple imperfect but effective results. Specifically, we prepare different prompts by sampling distinct examples for each prompt to utilize the in-context learning of LLM, and aggregate the results generated by each prompt with a Majority Vote strategy.

4 SEARCH SPACE OPTIMIZER

4.1 Dimensionality Optimization

In this part, *Search Space Optimizer* prunes space dimensionality by identifying knobs that have a significant impact on the DBMS performance, and only tunes these important knobs. Specifically, it utilizes LLM to simulate DBA’s empirical judgement in real-world scenarios and select knobs by considering the following four factors: (1) *DBMS Product.* After years of tuning practice, it is empirically known which knobs are important for a certain DBMS product. Since such wisdom is included in the corpus of GPT-4, we can extract it by prompting GPT-4 to recommend knobs based on the DBMS product. (2) *Workload Characteristics.* Different workloads have distinct requirements on DBMS resources, which are regulated by knobs. For example, given an I/O-intensive OLTP workload, DBAs would consider tuning disk-related knobs. (3) *Query Bottleneck.* Given LLM’s powerful analysis ability, LLM is capable of delving into the execution plans of queries, and choosing knobs related to the performance bottleneck. (4) *Knob Dependency.* Some knobs need to be tuned at the same time to take effect, and LLM can capture such dependency by reading DBMS manuals.

4.2 Range Optimization

In this part, *Search Space Optimizer* utilizes the structured knowledge in Section 3.2 to optimize the value range of each knob.

Region Discard. For numerical knobs, *Search Space Optimizer* limits the value scope of a knob between *min_value* and *max_value*. DBAs summarized such values in manuals, and we extract them to discard meaningless value regions, including values that are

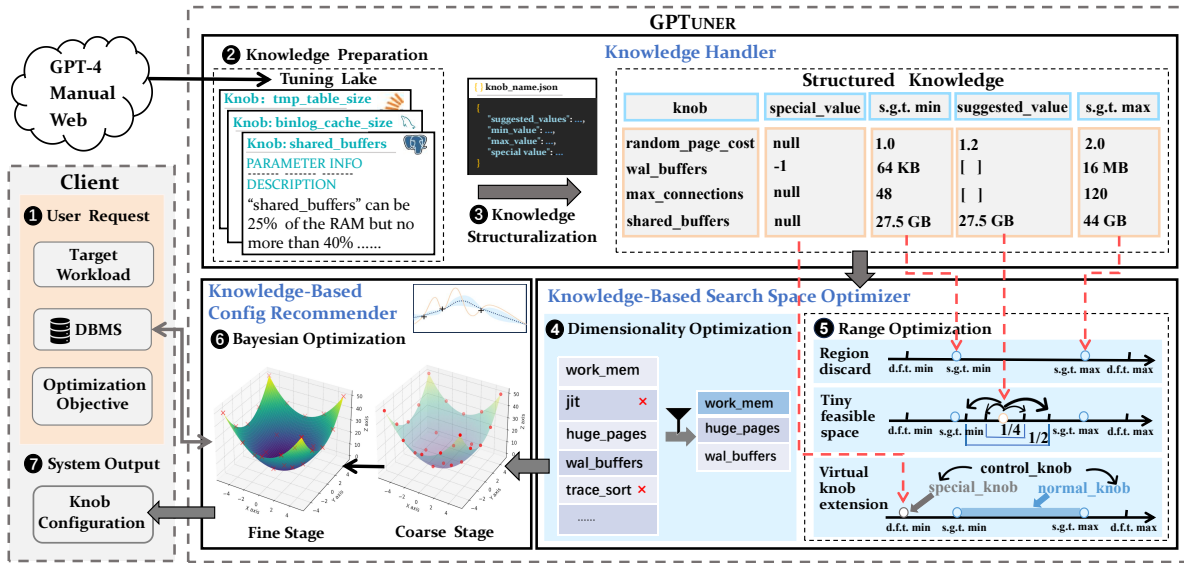


Figure 1: System Overview of GPTUNER

unlikely to result in promising performance, that could seize too many system resources, and that could even make the DBMS crash.

Tiny Feasible Space. *Search Space Optimizer* uses *suggested_values* to define a reliable discrete space. Such values are valuable since they performed well in the past and can serve as good starting points for the new scenario. However, they may not be suitable for all cases, as the optimal knob setting depends on the specific environment, which is diverse. Instead of relying on these static values only, we dynamically apply a set of multipliers on each suggested value. We conduct above deviation process for all numerical knobs, and the resulting discrete space is our *Tiny Feasible Space*, where *Tiny* means the possible number of values is significantly reduced, and *Feasible* indicates the chosen values are promising.

Virtual Knob Extension. In modern DBMS, there are knobs using special values to do something different from what the knobs normally do. However, optimizers may never trial these values (even though it could be the optimal) since the likelihood of sampling them is extremely low [2]. Therefore, we provide separate boolean virtual knobs to control each knob with special values. Specifically, the virtual knobs are exposed to optimizers to determine whether the “normal value range” or “special value” is activated.

5 CONFIGURATION RECOMMENDER

Configuration Recommender utilizes a novel Coarse-to-Fine BO Framework to explore the search space under the guidance of domain knowledge, and recommends well-performing configurations. In the first stage, BO only explores a discrete subspace of the whole heterogeneous space (the *Tiny Feasible Space* defined in Section 4.2). This subspace is small in size but promising to contain good configurations since we generate it based on the reliable domain knowledge. In the next stage, to avoid the overlooking problem of coarse-grained search (it is inevitable to lose some useful configurations for any space reduction technique), BO explores the heterogeneous space thoroughly with the other two optimizations in Section 4.2. Moreover, we use the samples from the first stage to bootstrap the surrogate model of BO in the second stage. After the

two stages, the recommender outputs the best-performing knob configurations found within the budget limits specified by users.

6 DEMONSTRATION SCENARIOS

This section demonstrates how visitors interact with GPTUNER to tune their DBMS in two scenarios. Visitors have the opportunity to invoke GPT models. However, since this takes time and money, we provide intermediate results in advance for better interactivity.

Database Administrators (DBAs). Let us now walk through a typical scenario where a DBA aims to optimize a frequently running workload in the company. She first provides basic information of the target DBMS in **A** and OpenAI API key can be optionally offered for model invocation. Then she uploads the workload in **B** and describes the characteristics of the workload for better knob selection. Clicking the button in **C**, GPTUNER identifies important knobs and DBA can check and modify the selection, and the final selection will be displayed in **D**. When she selects one of the knobs, the page updates to present the knowledge process pipeline for that knob. She can use knowledge from multiple sources, and provide her unique tuning insight as “Additional Suggestion”. Clicking the button in **F**, GPTUNER filters noisy knowledge contradicting the system view in **E**, and the remaining consistent knowledge will be summarized in **G**. Next, GPTUNER extracts the values of attributes from the summary, and generates the structured knowledge in **H**. At the same time, the optimized search spaces of the two stages are calculated and visualized in **I**. Finally, when she clicks “Start Tuning” in **J**, GPTUNER executes BO to explore the search spaces from coarse granularity to fine granularity. The tuning process is visualized and the best performance is reported in real time. All configurations are presented in **K**, and she can export the satisfied configuration in SQL format and deploy it on her DBMS directly.

Regular Engineers. Let us now walk through the second scenario where a regular engineer wants to try out GPTUNER for automatic DBMS optimization. In **B**, she can choose TPC-C or TPC-H as the target workload, which are supported by Benchmark [1]. She can just leave knob selection default and learn about the tuning knowledge of each knob in **E** and **F**. She can see how important attributes

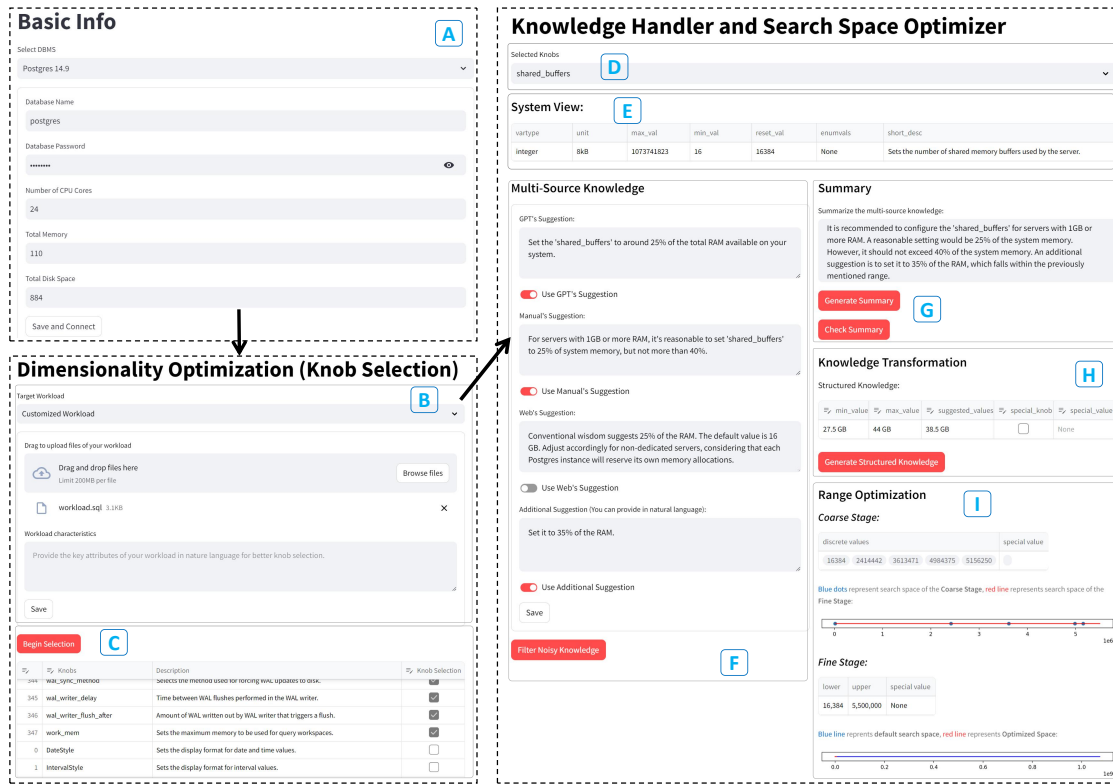


Figure 2: GPTuner Demonstration

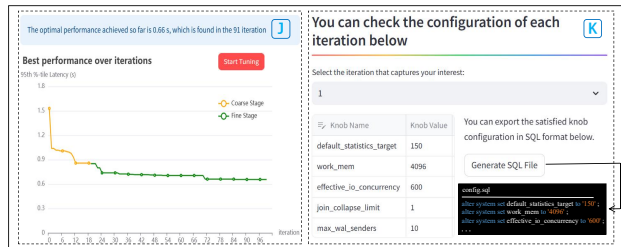


Figure 3: GPTuner Result Page

can be extracted from [G](#) to obtain the structured knowledge in [H](#), and see how efficiently GPTuner can reduce the search space provided by DBMS vendors, the red line and blue dots denote the search spaces of our two-stage BO algorithm. Finally, by clicking “Start Tuning” in [J](#), her target DBMS is optimized by GPTuner automatically.

ACKNOWLEDGEMENTS

Jianguo Wang acknowledges the support of the National Science Foundation under Grant Number 2337806.

REFERENCES

- [1] Djelle Eddine Difallah, Andrew Pavlo, Carlo Curino, and Philippe Cudré-Mauroux. 2013. OLTP-Bench: An Extensible Testbed for Benchmarking Relational Databases. *PVLDB* 7 (2013), 277–288.
- [2] Konstantinos Kanellis, Cong Ding, Brian Kroth, Andreas Müller, Carlo Curino, and Shivaram Venkataraman. 2022. LlamaTune: Sample-Efficient DBMS Configuration Tuning. *Proc. VLDB Endow.* 15 (2022), 2953–2965.
- [3] Jiale Lao, Yibo Wang, Yufei Li, Jianping Wang, Yunjia Zhang, Zhiyuan Cheng, Wanghu Chen, Mingjie Tang, and Jianguo Wang. 2023. GPTuner: A Manual-Reading Database Tuning System via GPT-Guided Bayesian Optimization. arXiv:2311.03157 [cs.DB]
- [4] Immanuel Trummer. 2021. The Case for NLP-Enhanced Database Tuning: Towards Tuning Tools That “Read the Manual”. *Proc. VLDB Endow.* 14, 7 (2021), 1159–1165.
- [5] Immanuel Trummer. 2022. DB-BERT: A Database Tuning Tool That “Reads the Manual” (*SIGMOD ’22*). Association for Computing Machinery, 190–203.
- [6] Bohan Zhang, Dana Van Aken, Justin Wang, Tao Dai, Shuli Jiang, Jacky Lao, Siyuan Sheng, Andrew Pavlo, and Geoffrey J. Gordon. 2018. A Demonstration of the Ottortune Automatic Database Management System Tuning Service. *Proc. VLDB Endow.* 11 (2018), 1910–1913.
- [7] Xinyi Zhang, Zhuo Chang, Yang Li, Hong Wu, Jian Tan, Feifei Li, and Bin Cui. 2022. Facilitating Database Tuning with Hyper-Parameter Optimization: A Comprehensive Experimental Evaluation. *Proc. VLDB Endow.* 15 (may 2022), 1808–1821.