Sawmill: From Logs to Causal Diagnosis of Large Systems

Markos Markakis MIT CSAIL Cambridge, MA USA markakis@mit.edu

Rana Shahout Harvard University Cambridge, MA USA rana@seas.harvard.edu An Bo Chen MIT CSAIL Cambridge, MA USA anbochen@mit.edu

Peter Baile Chen MIT CSAIL Cambridge, MA USA peterbc@mit.edu

Brit Youngmann Technion Haifa, Israel brity@technion.ac.il

Chunwei Liu MIT CSAIL Cambridge, MA USA chunwei@mit.edu Trinity Gao MIT CSAIL Cambridge, MA USA trinityg@mit.edu

Ibrahim Sabek University of Southern California Los Angeles, CA USA sabek@usc.edu Ziyu Zhang MIT CSAIL Cambridge, MA USA sylziyuz@mit.edu

Michael Cafarella MIT CSAIL Cambridge, MA USA michjc@csail.mit.edu

ABSTRACT

Causal analysis is an essential lens for understanding complex system dynamics in domains as varied as medicine, economics and law. Computer systems are often similarly complex, but much of the information about them is only available in long, messy, semistructured log files. This demo presents Sawmill, an open-source system [10] that makes it possible to extract causal conclusions from log files. Sawmill employs methods drawn from the areas of data transformation, cleaning, and extraction in order to transform logs into a representation amenable to causal analysis. It gives log-derived variables human-understandable names and distills the information present in a log file around a user's chosen causal units (e.g. users or machines), generating appropriate aggregated variables for each causal unit. It then leverages original algorithms to efficiently use this representation for the novel process of Explorationbased Causal Discovery - the task of constructing a sufficient causal model of the system from available data. Users can engage with this process via an interactive interface, ultimately making causal inference possible using off-the-shelf tools. SIGMOD'24 participants will be able to use Sawmill to efficiently answer causal questions about logs. We will guide attendees through the process of quantifying the impact of parameter tuning on query latency using real-world PostgreSOL server logs, before letting them test Sawmill on additional logs with known causal effects but varying difficulty. A companion video for this submission is available online [9].

CCS CONCEPTS

• Software and its engineering → System administration; • Computing methodologies → Causal reasoning and diagnostics; Natural language generation.

KEYWORDS

Logs, Fault Diagnosis, Causality, LLMs, Causal Discovery

SIGMOD-Companion '24, June 9–15, 2024, Santiago, AA, Chile © 2024 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-0422-2/24/06

https://doi.org/10.1145/3626246.3654731

ACM Reference Format:

Markos Markakis, An Bo Chen, Brit Youngmann, Trinity Gao, Ziyu Zhang, Rana Shahout, Peter Baile Chen, Chunwei Liu, Ibrahim Sabek, and Michael Cafarella. 2024. Sawmill: From Logs to Causal Diagnosis of Large Systems. In Companion of the 2024 International Conference on Management of Data (SIGMOD-Companion '24), June 9–15, 2024, Santiago, AA, Chile. ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3626246.3654731

1 INTRODUCTION

Failures are frequent in today's large, complex computer systems, and diagnosing them in production can be challenging [5]: there is usually not enough time (or often even the access permissions) for debugging techniques like testing [2], formal verification [4] and simulation [8]. Instead, operators have to work backward from failures using observational data collected from the system. Informally, we have heard of operations teams spending tens of hours with tools like Datadog [3], trying to diagnose a problem.

However, operations teams want to go beyond a diagnosis - they want to repair the system by alerting the appropriate engineering team. Moreover, whenever there are *multiple ways* to fix a problem, they would like to identify the *most efficient* way to utilize engineering effort. This points to the growing field of causal reasoning [13], which has provided scientists with a common language to express and evaluate hypotheses across diverse domains.

We aim for a general-purpose causal diagnosis tool that can help address a wide range of software systems problems by letting engineers *pose and correctly answer Average Treatment Effect (ATE) queries* [13]. Intuitively, the ATE formalizes a "dose-response" model under Pearl's theory of causality [13] (which we adopt in this work): for example, per "dose" of paralellism (e.g., one more worker), by how much will query latency decrease? Crucially, calculating ATEs correctly from observational data involves adjusting for *confounders*: common causes of both variables involved in an ATE calculation that could bias the observed effect [13]. For example, the available memory could impact both the chosen degree of parallelization and the query latency directly. Calculating ATEs quantifies the tradeoffs involved when several interventions can have *some* impact, helping engineers pick the most efficient way forward.

Unfortunately, directly calculating ATEs based on the raw data available to large systems operators is impossible. The raw data usually take the form of collections of logs: semi-structured chronological accounts in text form. Figure 1a presents a snippet of a

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).







real-world log from the POSTGRESQL dataset that we will use in our demonstration. Past work on log management has led to significant infrastructure, including tools offering extensive coverage of software and hardware components [7]. However, causal inference toolkits require a different representation of the underlying data [17], like that of Figure 1b – a table with *one row per data point*, including only *few*, *relevant variables* and *without missing data*. They also require a causal model [13], something non-trivial to recover from a log. A causal model captures domain knowledge about the relationships among variables and is often represented as a directed acyclic graph (DAG) [13], like that of Figure 1c: each node represents a variable and each directed edge encodes a direct influence of the source variable on the destination variable.

In summary, we aim to combine log management with the best theoretical machinery causality can offer, but face three challenges:

Challenge A: Deriving the Schema. Text logs are a far cry from the tabular dataset in Figure 1b. Log parsing algorithms can convert log data into *some* tabular representation, but can yield hundreds of variables without a human-friendly way to navigate them, like the interpretable column names in Figure 1b.

Challenge B: Distilling the Data. Causal inference requires the same data per causal unit - e.g., per session. Simply parsing the log falls short of this goal. First, there will be a lot of missing values because each log message only reports *some* variables. Second, logs record information at a very fine granularity. This information must be summarized, while preserving "causal usefulness".

Challenge C: Obtaining the Causal Model. For every possible pair of variables, an expert could easily decide the plausibility and direction of the causal relationship between them. However, fully specifying a model over all the log variables by hand is daunting given the problem size. Another option could be to derive the model from the data, a process called causal discovery [13]. However, existing algorithms would fall short due to the problem size and



Figure 2: Our framework, indicating the three challenges.

the functional dependencies among variables, while large language models would be tripped up by the context-specific variables.

We propose a framework to address these challenges and transform logs into high-grade causal resources to enable rapid diagnosis of system failures. We implemented this framework in Sawmill, a system that helps engineers managing complex systems formulate a data-driven hypothesis about the cause of an observed failure and retrieve the correctly adjusted associated Average Treatment Effect (ATE). To address **Challenge A**, we derive human-understandable variable tags by leveraging Large Language Models; for **Challenge B**, we distill log information around *causal units* and generate new variables for every causal unit to maximize "usefulness"; while in the face of **Challenge C**, we propose interfaces to recover a relevant, partial causal model of the system from log-derived data.

SIGMOD '24 participants will use Sawmill to investigate the causes of system failures starting from system logs. This includes inspecting the raw log files, examining the generated dataset, and finding the primary causes of unexpected behavior. We will walk participants through an analysis of *real-world* POSTGRESQL logs, where max_parallel_workers and work_mem confound each other's impact on downstream latency on TPC-DS [15] (green edge in Figure 1c). We will then give them the chance to independently use Sawmill to analyze failures in more log datasets of varying difficulty.

2 SAWMILL ARCHITECTURE

Figure 2 summarizes our framework, which transforms logs into a table like Figure 1b and derives a causal graph like Figure 1c, enabling ATE calculations. It starts by **Parsing** logs and **Tagging** each variable with a human-understandable tag. The parsed information is then reorganized to satisfy the requirements for causal inference, by **Defining Causal Units** and **Computing Suitable Variables**, making **Obtaining a Causal Graph** and **Answering ATE Queries** possible. We will now dive deeper into each step.

2.1 From the Log to the Parsed Table

2.1.1 Log Parsing. Using the function PARSE, logs are first parsed into the *parsed table*, with one row per log message and one column per *parsed variable*. Users can optionally pass into PARSE regular expressions for variables like timestamps. PARSE then uses Drain [6], an off-the-shelf single-pass log parsing algorithm, to extract the rest of the parsed variables. We discard any categorical variable with over 0.15V distinct values across V occurrences, similar to past work [19]. If Drain has mapped different log variables to the same parsed variable, the user can identify it in the parsed table and correct it using Sawmill's SEPARATE function.

Markakis et al.

2.1.2 Variable Tagging. For variables parsed using a regular expression, the user provides a tag together with the regular expression. For the rest, Sawmill automatically generates a unique humanunderstandable tag by consulting three sources in sequence: the tokens preceding the variable in the corresponding *log template*; GPT-3.5-Turbo [12], given an example message including the variable and example values for the variable from other log messages; and GPT-4 [11], given the same information. If no source produces a tag, or if the produced tag is already assigned to a different variable, Sawmill assigns a unique string of symbols instead.

2.2 From the Parsed Table to the Prepared Table

2.2.1 Defining Causal Units. The parsed table includes one row per log message, but useful ATE queries for troubleshooting are usually about larger units, like sessions or machines. This requires grouping several log messages together as a *causal unit*, with each group being "one data point" for the ATE query at hand. To define causal units in Sawmill, a user calls SETCAUSALUNIT on the appropriate parsed variable - e.g. sessionID. Not every choice of causal unit is permissible, because of the Stable Unit Treatment Value Assumption (SUTVA) [16]: the outcome of each unit should not depend on the treatments of *other units*. For example, users may be unsuitable causal units if they share hardware: user *A*'s work could impact user *B*'s latency. We defer to the user's knowledge of the particular system at hand to ensure that this condition is satisfied.

2.2.2 Computing Suitable Variables. There can be a varying number of values for each parsed variable in each causal unit. For example, each may have a different number of query latency readings. To make units comparable, Sawmill replaces such varying-size collections of values with the same aggregate(s) of the values per causal unit. This transformation, triggered using PREPARE, yields the prepared table, with one row per causal unit and one column for each prepared variable. Each prepared variable is derived from some parsed variable (its base variable) by using some function (e.g. the mean) to reconcile the values within each causal unit. A number of prepared variables are generated for each parsed variable using different functions. Among prepared variables sharing a base variable, Sawmill then only keeps the one that maximizes empirical entropy, to limit the prepared table size and processing cost. We also let the user optionally impute missing values with a default based on domain knowledge, by passing the value to PREPARE.

2.3 From the Prepared Table to ATEs

2.3.1 Obtaining a Causal Graph. To obtain a causal graph, Sawmill combines the data-driven and expert-driven approaches: it helps users incrementally build the relevant part of the causal graph, by making data-driven suggestions that the user evaluates using expert knowledge. We call this approach *Exploration-based Causal Discovery*. The user begins with an "outcome" variable *Y* (e.g mean query latency) and builds the causal graph around it in two phases:

• **Phase I - Finding a "root cause":** Sawmill helps identify a treatment prepared variable *T*, which *causally affects Y* and is also "actionable". It does so by providing the function EXPLORECANDIDATECAUSES(*U*), which suggests likely

causes for a prepared variable U. By iteratively leveraging it, the user can arrive at T, concluding Phase I. To efficiently leverage the user's attention, Sawmill only presents a *pruned* list of candidate causes, obtained using LASSO [18]. Because each variable is linearly related to its parents and the set of parents is often expected to be small [1], variables that get assigned a zero coefficient can be safely ignored. The user can inspect each candidate cause U', ranked by increasing p-value, and decide whether to include $U' \rightarrow U$ in the causal graph (ACCEPT(U', U)) or not (REJECT(U', U)).

• Phase II - Finding confounders: The user continues constructing the causal graph to identify a sufficient set of confounders that affect ATE(T, Y). We measure the user's progress in recovering the regions of the causal graph that *could* include confounders by calculating an *exploration score*: the fraction of accepted/rejected edges among edges that touch at least one node in the current graph. Phase II ends when the exploration score reaches 1. We provide SUGGEST-NEXTEXPLORATION, which returns a prepared variable U_s such that calling EXPLORECANDIDATECAUSES(U_s) will yield as many edges relevant to the exploration score calculation as possible, maximizing the user's decision-making efficiency.

2.3.2 Answering ATE Queries. Having transformed the log data into the prepared table and having obtained the relevant part of the causal graph, the user can now calculate the ATE of interest using the function GETATE(T, Y).

3 PROTOTYPE AND DEMONSTRATION

3.1 Prototype System

Sawmill uses ~2800 lines of Python and is open-source [10]. We used "backdoor.linear_regression" from DoWhy [17], a Python causal inference library, for GETATE and scikit-learn [14] for LASSO.

3.2 Guided Demonstration

In this part of our demo, SIGMOD attendees will act as operators at a database-as-a-service company. Customers use different configurations for the company's database, with some experiencing higher mean latency for the same workload. Attendees are interested in correctly determining the impact of the max_parallel_workers parameter on mean latency. They have access to the **PostGRESQL** dataset: logs collected on PostgreSQL 14, into which we loaded TPC-DS [15] for scale factor 1 and sequentially issued the TPC-DS queries, excluding four long-running queries. We ran this workload for different settings of the parameters from Figure 1c.

Step 1: Obtaining the Parsed Table Users inspect the log and invoke PARSE, as shown in Figure 3a. Inspecting the parsed table reveals a parsing miss: distinct variables have been mapped to a single parsed variable. Users call SEPARATE on the mis-parsed variable to instruct Sawmill to correct the parsing mistake.

Step 2: Obtaining the Prepared Table Users set each session as a causal unit using SetCAUSALUNIT, before invoking PREPARE.

Step 3: Calculating the ATE Users ask Sawmill for candidate causes of mean query latency using EXPLORECANDIDATECAUSES, yielding the results of Figure 3b. These candidates include 6 variables related to the modified PostrgeSQL parameters, so users include the corresponding edges in the graph using ACCEPT. Users

SIGMOD-Companion '24, June 9-15, 2024, Santiago, AA, Chile

Parsing and Tagging											
We can call the parse() function to parse tree separate each log line into <i>templates</i> it. Sawmill then assigns a human-underst	e the log into a table using the and a collection of <i>log variab</i> andable <i>tog</i> to each parsed va	e <u>Drain</u> algorithr les, based on th ariable.	n. The Drain algorithm utilizes a e similarity of each new line to	fixed-depth parse the lines seen before							
Force re-calculation											
Message Prefix	0	Custom Regex									
$\ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ \ $	The Drain algorithm can first extract named variables from each log										
Similarity Threshold	0	line using regular expressions, which the users can provide here:									
	0.65	-	≂, value								
0.00	1.00	Date	\d{4}-\d{2}-\d{2}								
Depth	0	Time	\d{2}:\d{2}:\d{2}\.\d{3}(?=EST \[}								
5	- +	sessionID	(?<=EST \[)\S+\.\S+								
		tID	3/\d+(?=])								
Parse											
Paise											

(a) Users can edit the arguments to PARSE, among other functions.

Explore candidate causes

1									
Choose a variable to	Candida	Candidate Cause(s)							
explore candidate causes for: Select available duration mean		Candidate	Candidate Tag	Target	Slope	P-value	Candidate->1		
	0	5da61167_14+mean	time mean	9989ce8d_12+mean	3.3003	0	Undecided		
	1	da4032f7_15+mean	work_mem mean	9989ce8d_12+mean	-123.4389	0	Undecided		
	2	Date+mean	Date mean	9989ce8d_12+mean	-0.0393	0	Undecided		
	3	8d5769b2_15+mean	seq_page_cost mean	9989ce8d_12+mean	11.1776	0	Undecided		
Explore Candidate Causes Candidate cause(s) found1	4	9d2f8f03_15+mean	max_parallel_workers mean	9989ce8d_12+mean	5,127.6471	0.0122	Undecided		
	5	Time+mean	Time mean	9989ce8d_12+mean	0.0703	0.11	Undecided		
	6	9eea70b1_15+mean	random_page_cost mean	9989ce8d_12+mean	-0.9126	0.662	Undecided		
	7	aea309d3_23+mean	port mean	9989ce8d_12+mean	0.0244	0.8568	Undecided		
	8	5da61167_26+mean	5da61167_26 mean	9989ce8d_12+mean	0.0244	0.8568	Undecided		
	9	97a1388a_15+mean	maintenance_work_memme	9989ce8d_12+mean	0.0012	0.9854	Undecided		

(b) An invocation of ExploreCandidateCauses.



(c) Users can build a causal graph and monitor their ATE query.

Figure 3: Indicative snapshots of using Sawmill.

then invoke GETATE, but receive a perplexing result of 374.47 - **more parallelism means a longer duration**! Something is amiss.

Step 4: Adjusting for Confounding Users follow the output of Sawmill's SUGGESTNEXTEXPLORATION and explore candidate causes for max_parallel_workers:mean. They find that work_mem:mean is the top candidate - working memory confounds the effect of parallelism on latency! Users invoke ACCEPT to add the appropriate edge to the graph, recreating the graph of Figure 1c in Figure 3c. As shown on the left, this takes the ATE of interest to -156.47 - more parallelism yields a lower query duration, as expected.

3.3 Independent Exploration

SIGMOD attendees can use Sawmill on two more dataset families:

PROPRIETARY: This dataset is based on a real log for an HTTPbased application from a large company. It covers a collection of users, a fraction F of which is on a faulty OS version, failing HTTP requests with probability p_f . Attendees will examine Sawmill's ability to recover the ATE of the faulty OS version on HTTP failure responses, for different values of F and p_f .

XYZ: This dataset logs *V* synthetic variables for each of a set of "machines". While the values of most variables are randomly chosen, those of *x*, *y* and *z* are not. We have that ATE(x,y)=2, but this is distorted by confounding by *z* and by Gaussian noise with $\sigma = R$. Attendees will examine Sawmill's ability to reliably detect and adjust for *z*'s confounding, for different values of *V* and *R*.

ACKNOWLEDGMENTS

We gratefully acknowledge the support of the DARPA ASKEM project (Award No. HR00112220042).

REFERENCES

- Tom Claassen, Joris Mooij, and Tom Heskes. 2013. Learning Sparse Causal Models is not NP-hard. arXiv:1309.6824 [cs.AI]
- [2] Ermira Daka and Gordon Fraser. 2014. A Survey on Unit Testing Practices and Problems. In 2014 IEEE 25th International Symposium on Software Reliability Engineering. IEEE Computer Society CPS, Los Alamitos, CA, USA, 201–211. https: //doi.org/10.1109/ISSRE.2014.11
- [3] Datadog. [n. d.]. Datadog. Retrieved Apr 3, 2024 from https://www.datadoghq. com
- [4] Vijay D'silva, Daniel Kroening, and Georg Weissenbacher. 2008. A Survey of Automated Techniques for Formal Software Verification. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems 27, 7 (2008), 1165–1178.
- [5] Saurabh Gupta, Tirthak Patel, Christian Engelmann, and Devesh Tiwari. 2017. Failures in Large Scale Systems: Long-term Measurement, Analysis, and Implications. In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (Denver, Colorado) (SC '17). Association for Computing Machinery, New York, NY, USA, Article 44, 12 pages. https://doi.org/10.1145/3126908.3126937
- [6] Pinjia He, Jieming Zhu, Zibin Zheng, and Michael R. Lyu. 2017. Drain: An Online Log Parsing Approach with Fixed Depth Tree. In 2017 IEEE International Conference on Web Services (ICWS). IEEE Computer Society CPS, Los Alamitos, CA, USA, 33–40. https://doi.org/10.1109/ICWS.2017.13
- [7] Shilin He, Pinjia He, Zhuangbin Chen, Tianyi Yang, Yuxin Su, and Michael R Lyu. 2021. A Survey on Automated Log Analysis for Reliability Engineering. ACM computing surveys (CSUR) 54, 6 (2021), 1–37.
- [8] James Kapinski, Jyotirmoy V Deshmukh, Xiaoqing Jin, Hisahiro Ito, and Ken Butts. 2016. Simulation-based Approaches for Verification of Embedded Control Systems: An Overview of Traditional and Advanced Modeling, Testing, and Verification Techniques. *IEEE Control Systems Magazine* 36, 6 (2016), 45–64.
- [9] Markos Markakis. 2024. Sawmill Companion Video. Retrieved Mar 26, 2024 from https://youtu.be/uCS3pJZ45DI
- [10] Markos Markakis, An Bo Chen, Trinity Gao, and Peter Baile Chen. 2024. Sawmill Code. Retrieved Apr 3, 2024 from https://github.com/mitdbg/sawmill
- [11] OpenAI. 2023. GPT-4 Technical Report. arXiv:2303.08774 [cs.CL]
- [12] OpenAI. 2024. Models. Retrieved Mar 26, 2024 from https://platform.openai. com/docs/models/gpt-3-5
- [13] Judea Pearl. 2009. Causality: Models, Reasoning and Inference. Cambridge University Press, Cambridge, UK.
- [14] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [15] Meikel Poess, Bryan Smith, Lubor Kollar, and Paul Larson. 2002. TPC-DS, Taking Decision Support Benchmarking to the Next Level. In Proceedings of the 2002 ACM SIGMOD International Conference on Management of Data. Association for Computing Machinery, New York, NY, USA, 582–587. https://doi.org/10.1145/ 564691.564759
- [16] Donald B. Rubin. 1980. Comment on 'Randomization Analysis of Experimental Data in the Fisher Randomization Test'. J. Amer. Statist. Assoc. 75, 371 (1980), 591–93.
- [17] Amit Sharma, Emre Kiciman, et al. 2019. DoWhy: A Python Package for Causal Inference. https://github.com/microsoft/dowhy.
- [18] Robert Tibshirani. 1996. Regression Shrinkage and Selection via the Lasso. Journal of the Royal Statistical Society: Series B (Methodological) 58, 1 (1996), 267–288.
- [19] Wei Xu, Ling Huang, Armando Fox, David Patterson, and Michael I. Jordan. 2009. Detecting Large-Scale System Problems by Mining Console Logs. In Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles (Big Sky, Montana, USA) (SOSP '09). Association for Computing Machinery, New York, NY, USA, 117–132. https://doi.org/10.1145/1629575.1629587