



Adversarial Workload Matters

Executing a Large-Scale Poisoning Attack against Learned Index Structures

Matthias Bachfischer (Student ID: 1133751)

Melbourne, October 18th 2021

Agenda


- 1 Introduction**
- 2 Background on Adversarial ML**
- 3 Poisoning Attacks on Learned Index Structures**
- 4 Experimental Setup**
- 5 Experimental Results**
- 6 Conclusion**

Agenda

- 1** Introduction
- 2** Background on Adversarial ML
- 3** Poisoning Attacks on Learned Index Structures
- 4** Experimental Setup
- 5** Experimental Results
- 6** Conclusion

Traditional DB indexes use tree data structures to find record on disk – learned indexes uses ML models to “predict” location


Traditional database indexes



Type: B-Tree (and variantes)

Complexity: $O(\log_2 n)$

Learned Index Structures (LIS)




Type: Linear regression (and variantes)


Complexity: $O(1)$

A learned index structure works by approximating the CDF of the data to predict the location of the query key

Frequency Distribution

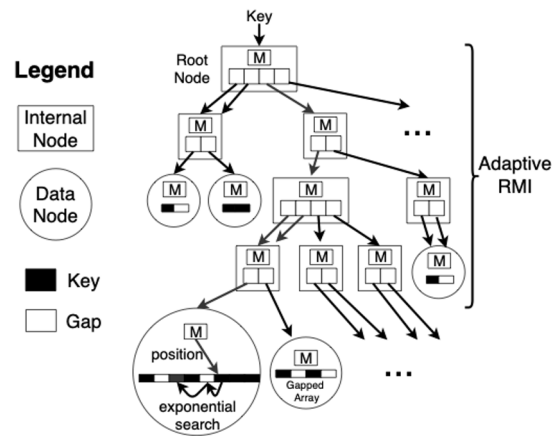


Cumulative Distribution Function (CDF)



Students
Amanda
Amy
Andrew
Andrew
Anna
Anna
Ashley
Barbara
Betty
Brenda
Brian
Carol
Carol
Christopher
Daniel
Daniel
David
Elizabeth
Emily
Eric
Gary
George
...
Susan
Thomas
William
Zamantha

This research has evaluated three index implementations: Two learned index structures (ALEX & D-PGM) and one B+Tree



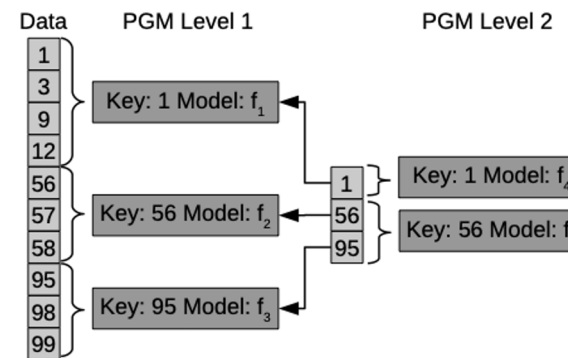
ALEX

Paper: ALEX: an updatable adaptive learned index

Authors: Ding et al. (2019)

Model: Piecewise Linear Approximation (PLA)

Code: <https://github.com/microsoft/ALEX>



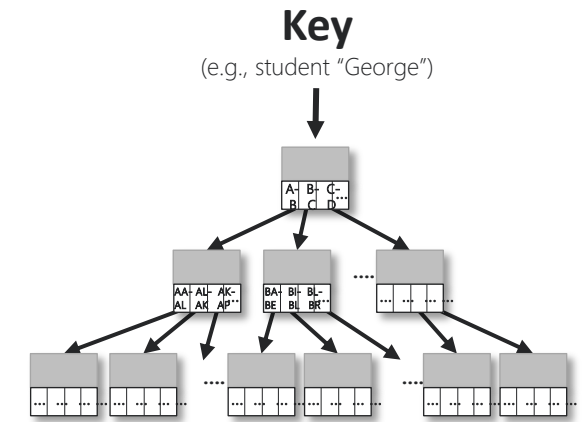
Dynamic-PGM

The PGM-index: a fully-dynamic compressed learned index with provable worst-case bounds

Ferragina et al. (2020)

Piecewise Linear Approximation (PLA)

<https://github.com/gvinciguerra/PGM-index>



B+Tree

STX B+ Tree C++ Template Classes v0.9

Timo Bingmann (2007)

None

<https://github.com/bingmann/stx-btree>

Agenda

- 1** Introduction
- 2** Background on Adversarial ML
- 3** Poisoning Attacks on Learned Index Structures
- 4** Experimental Setup
- 5** Experimental Results
- 6** Conclusion

Adversarial Machine Learning is the study of ML techniques against an adversarial opponents aiming to fool the model

		Attacker goals		
		Integrity	Availability	Privacy / Confidentiality
Attacker capability	Test data	<p>Misclassifications that do not compromise normal system operation</p> <p>Evasion (a.k.a. adversarial examples)</p>	<div></div>	<p>Examples that reveal secret information about the model or its users</p> <p>Model inversion (e.g., stealing model parameters)</p> <p>Model extraction / stealing</p> <p>Model inversion (e.g., stealing model parameters)</p> <p>Model inversion (e.g., stealing model parameters)</p>
	Training data	<p>Poisoning to allow subsequent intrusions</p> <p>(a.k.a. backdoors or neural network trojans)</p>	<p>Focus of this work</p> <p>Poisoning to maximize classification error</p> <p>(a.k.a. worst-case behavior)</p>	<p>n/a</p>

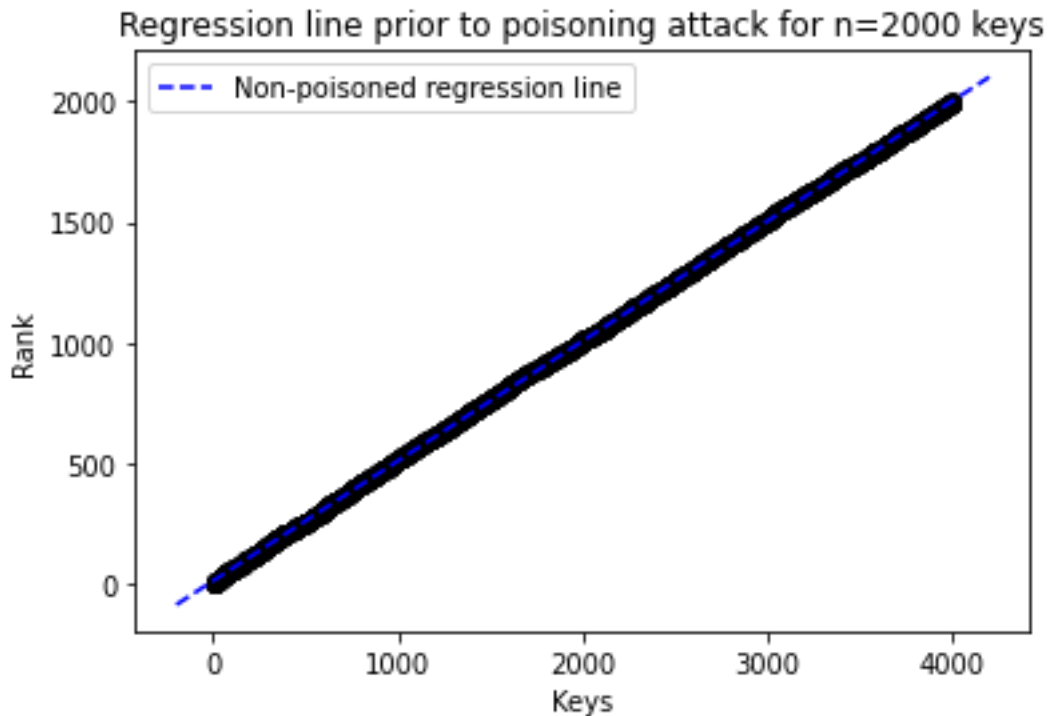
Source: Biggio, B. and F. Roli (2018). *Wild patterns: Ten years after the rise of adversarial machine learning*. Pattern Recognition 84: 317-331.

Agenda

- 1** Introduction
- 2** Background on Adversarial ML
- 3** Poisoning Attacks on Learned Index Structures
- 4** Experimental Setup
- 5** Experimental Results
- 6** Conclusion


To study the performance of learned indexes under adversarial workload, we execute a poisoning attack against the CDF

1. Dataset **before** poisoning attack:



Mean Squared Error (MSE): 22.99


2. Dataset **after** poisoning attack:






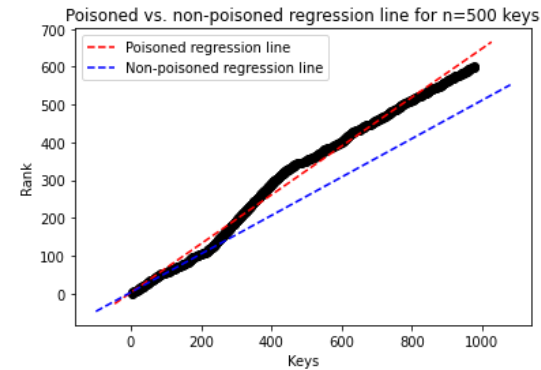
Mean Squared Error (MSE): 8675.83

Poisoning threshold $p = 20\%$ (400 keys)

Distributed poisoning algorithm deliberately places poisoning keys in dense areas to exacerbate the non-linearity of the CDF



Algorithm: GreedyPoisoningRegressionCDF¹
(based on Kornaropoulos et al.)



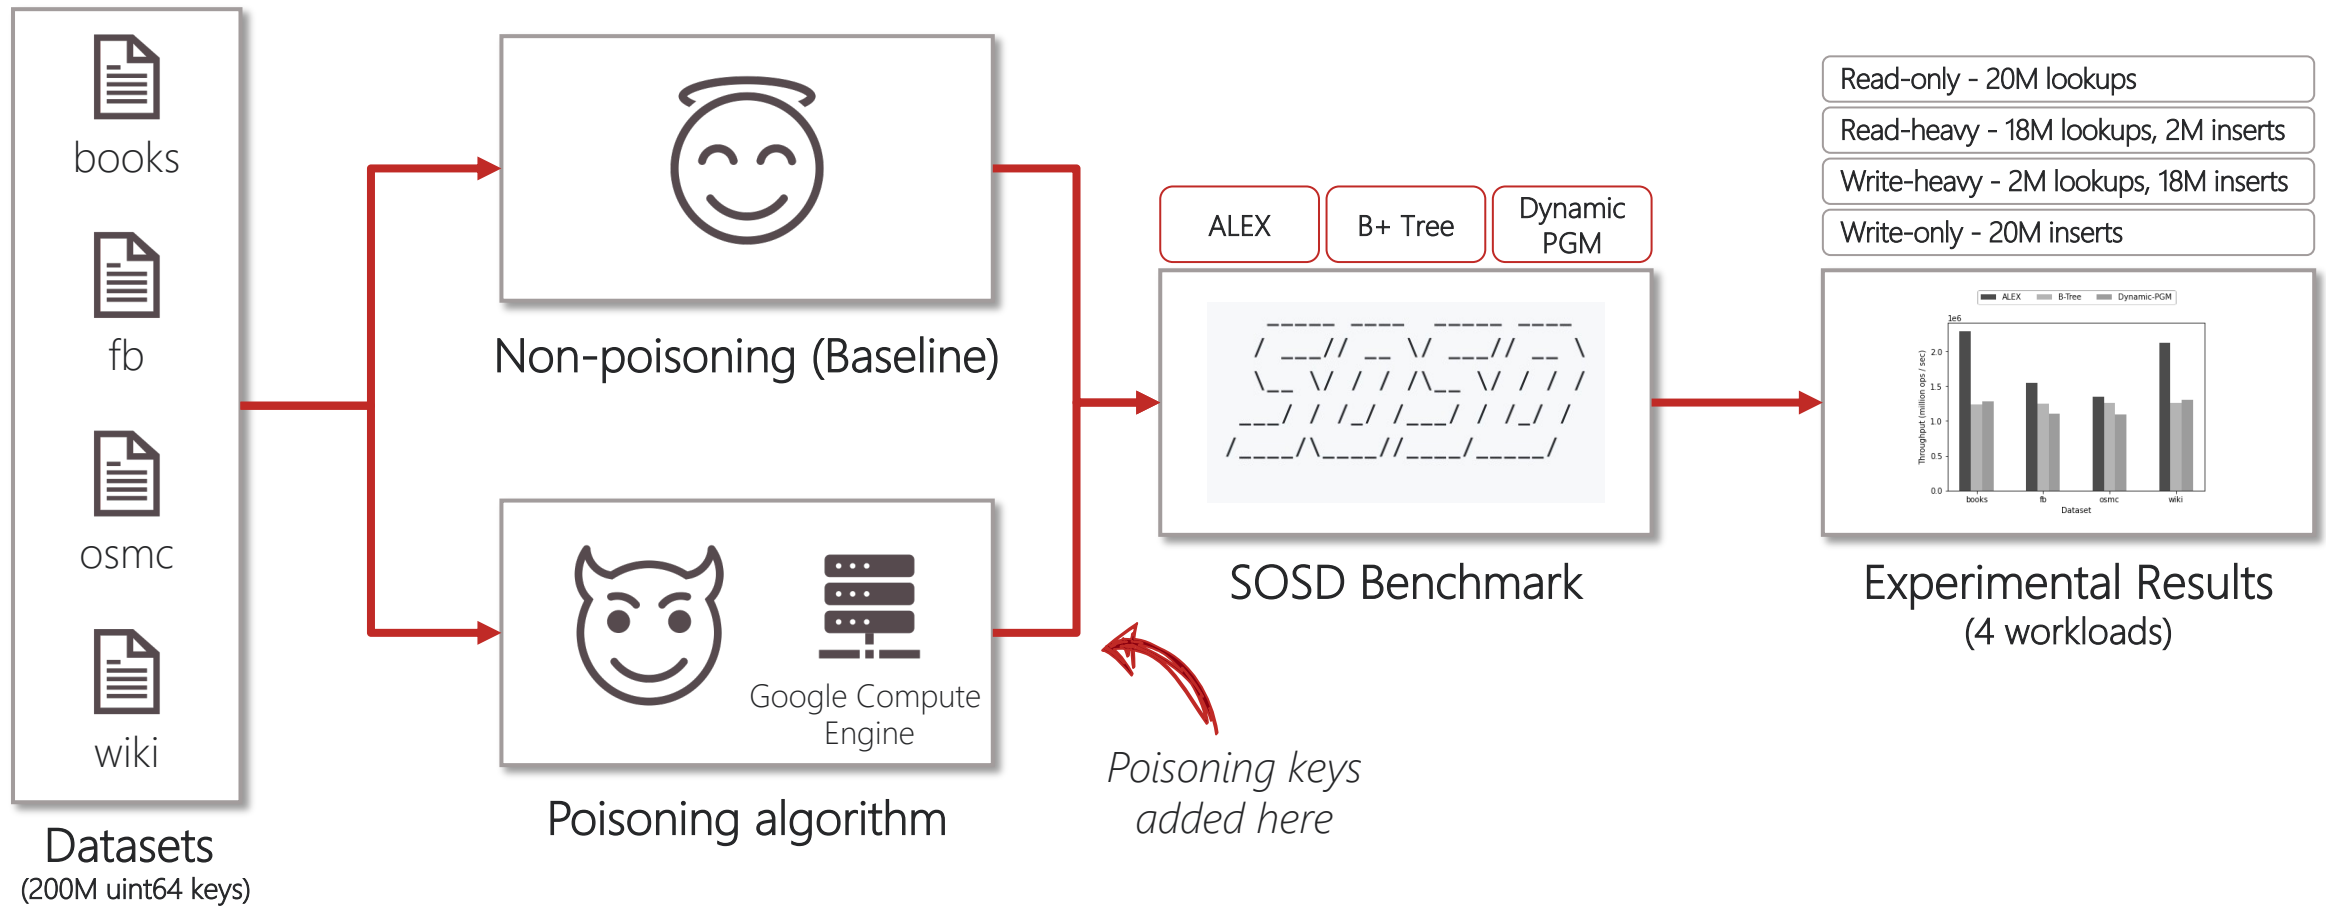
Algorithm: Greedy**Distributed**PoisoningRegressionCDF

1) Kornaropoulos, E. M., et al. (2020). *The Price of Tailoring the Index to Your Data: Poisoning Attacks on Learned Index Structures*.


Agenda

- 1** Introduction
- 2** Background on Adversarial ML
- 3** Poisoning Attacks on Learned Index Structures
- 4** Experimental Setup
- 5** Experimental Results
- 6** Conclusion


Data pipeline extends existing open-source tools from learned indexes research community




Key distribution of SOSD benchmark datasets varies heavily in terms of their cumulative distribution function




books_200M_uint64



fb_200M_uint64



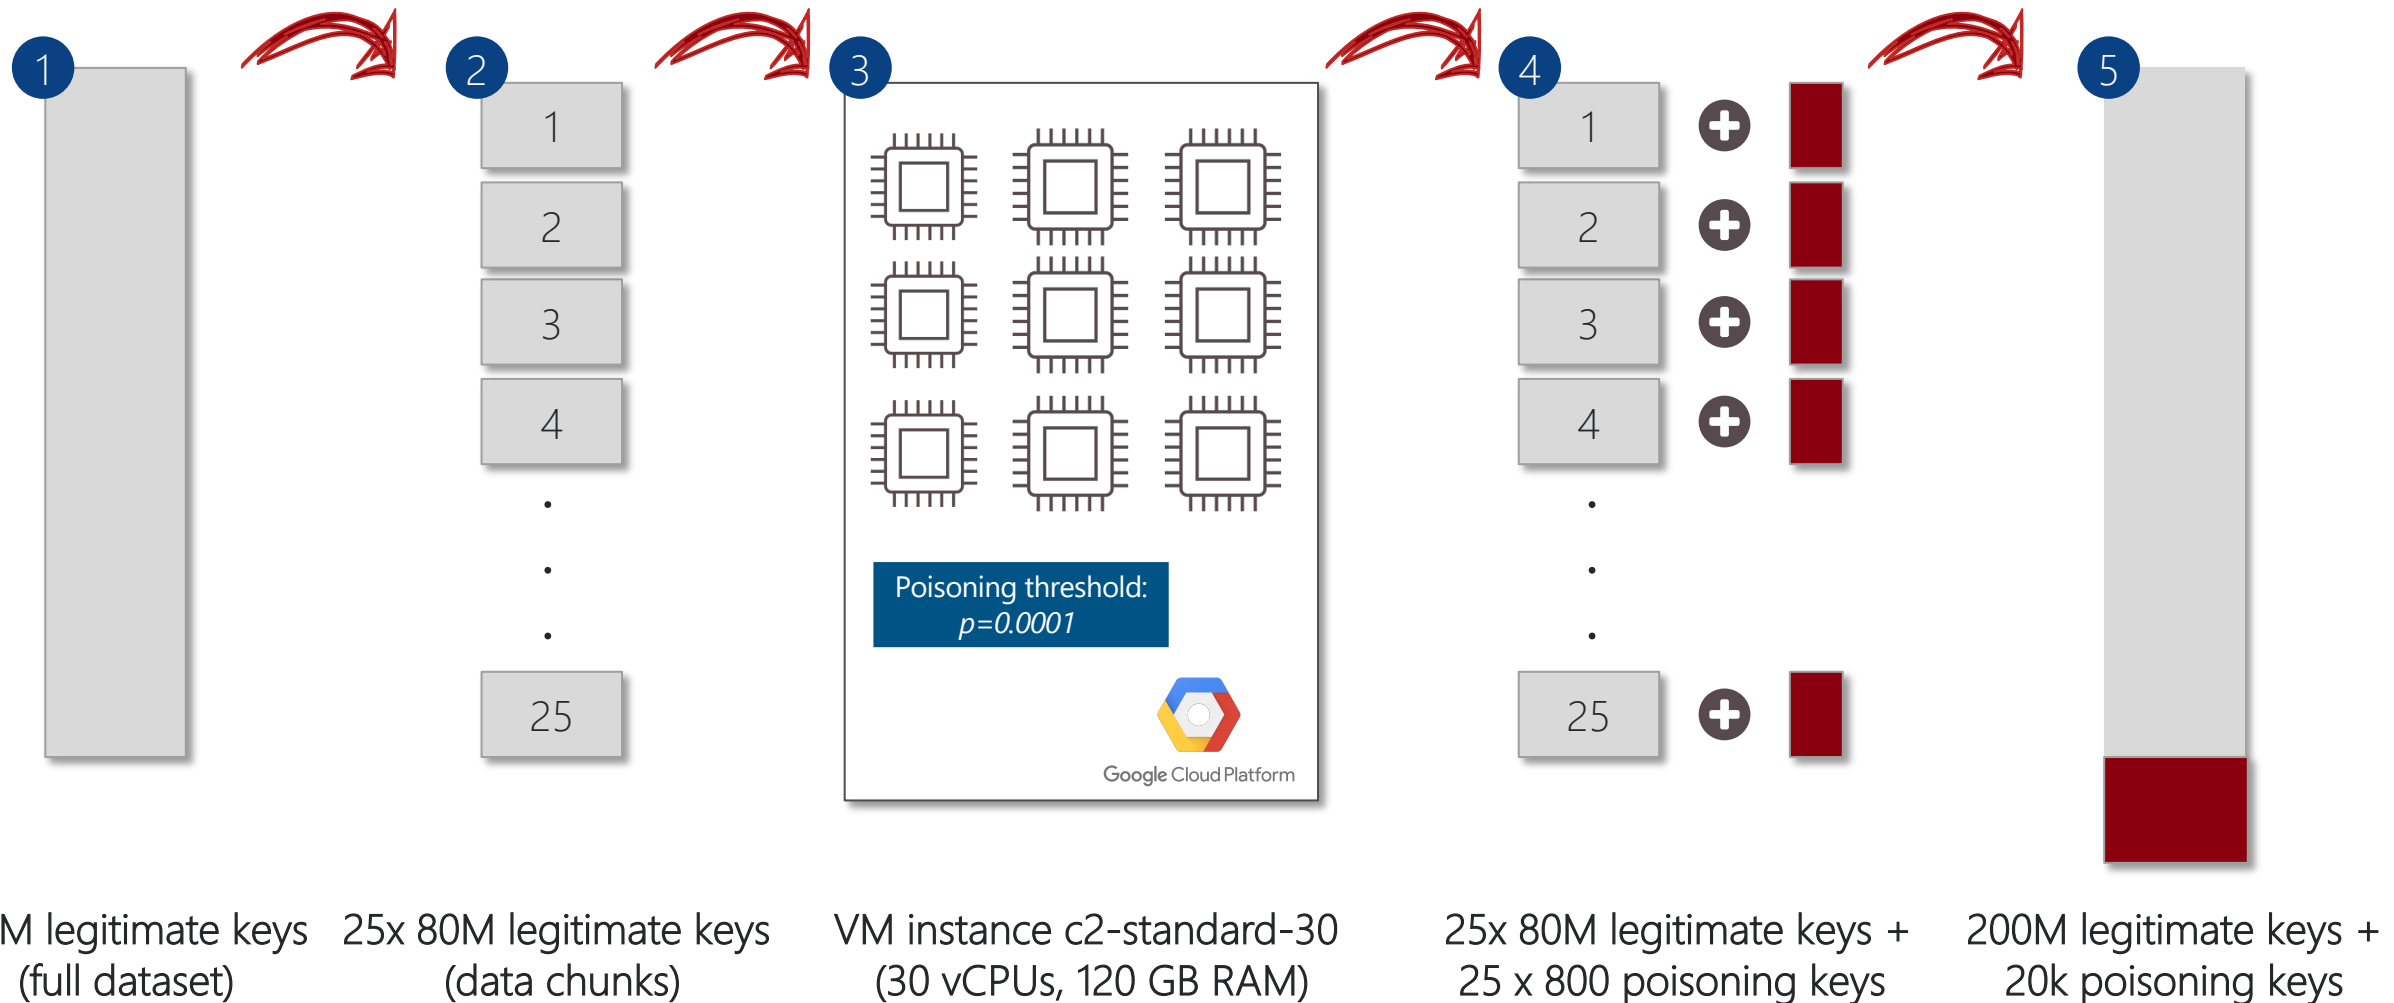
osm_cellids_200M_uint64



wiki_ts_200M_uint64

Contents:	Book popularity on Amazon	Facebook user IDs	Cell IDs on Open Street Map	Wikipedia edit timestamps
Data:	200M uint64 keys	200M uint64 keys	200M uint64 keys	200M uint64 keys
Size:	1.53 GB	1.53 GB	1.53 GB	1.53 GB

Poisoning attack was scaled to target up to 200M keys by exploiting the power of a very large virtual machine



Note: Poisoning attack was executed on a c2-standard-30 VM instance running on Google Compute Cloud (30 vCPUs, 120GB RAM).

Agenda

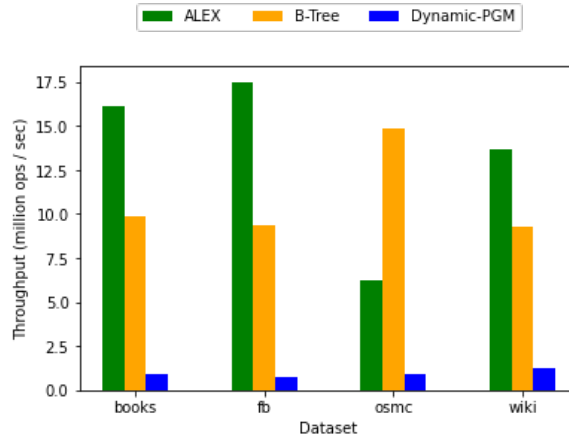
- 1** Introduction
- 2** Background on Adversarial ML
- 3** Poisoning Attacks on Learned Index Structures
- 4** Experimental Setup
- 5** Experimental Results
- 6** Conclusion

Final results from poisoned and non-poisoned workload scenarios with poisoning percentage $p=0.0001$ (I/II)

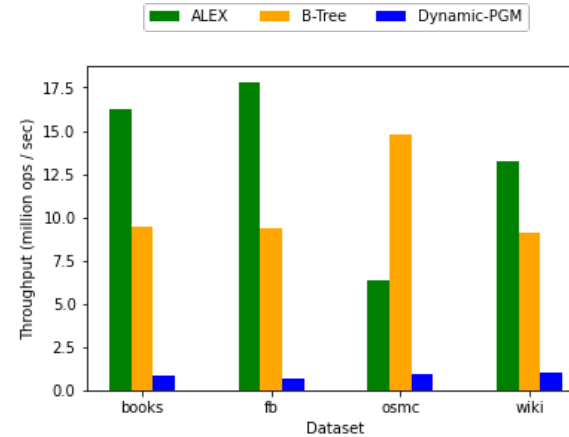
1. Read-only workload

20M lookups

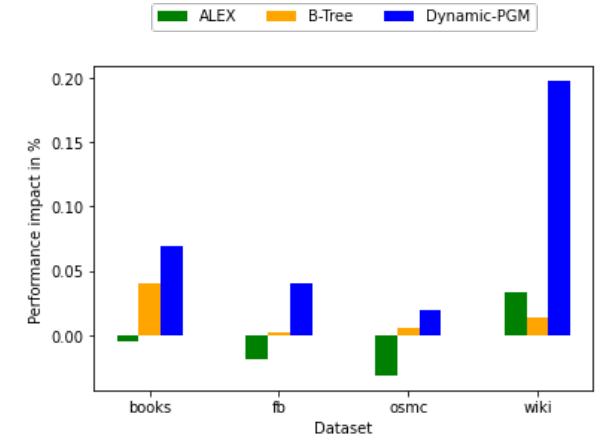
Non-poisoned scenario:



Poisoned scenario:


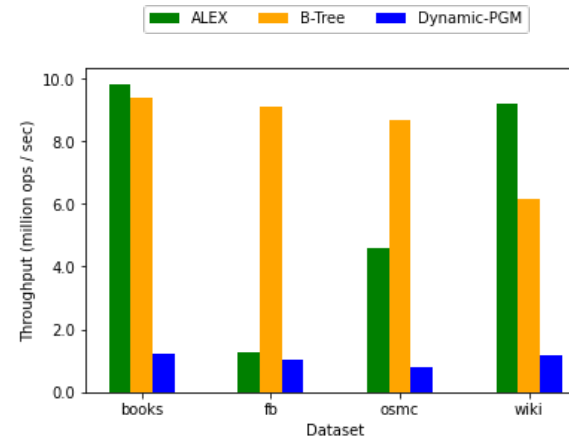
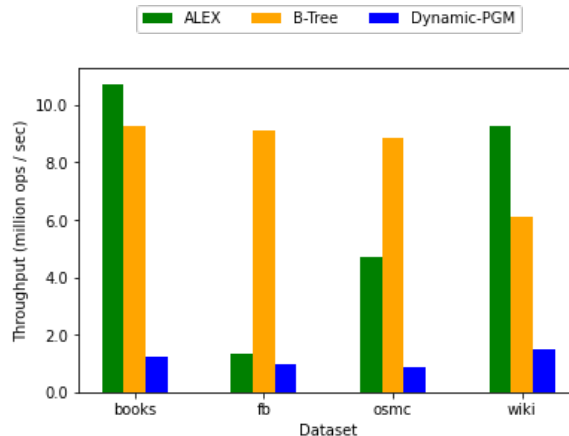


Performance impact:



2. Read-heavy workload

18M lookups,
2M inserts



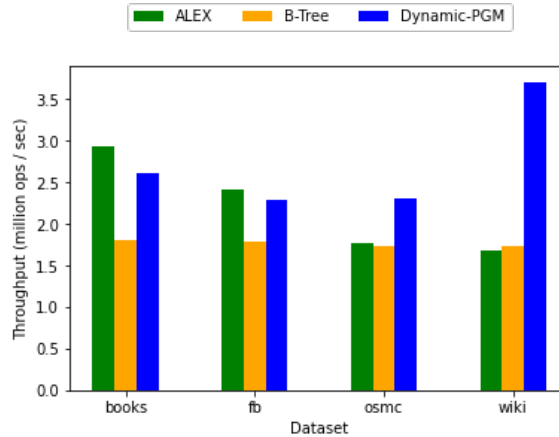
Note: All datasets in this experiment consist of 200M uint64 keys. The benchmark was evaluated on a *e2-standard-8* VM instance running on Google Compute Cloud (8 vCPUs, 32GB RAM). For each index, multiple configurations were tested..

Final results from poisoned and non-poisoned workload scenarios with poisoning percentage $p=0.0001$ (II/II)


3. Write-heavy workload

2M lookups,
18M inserts

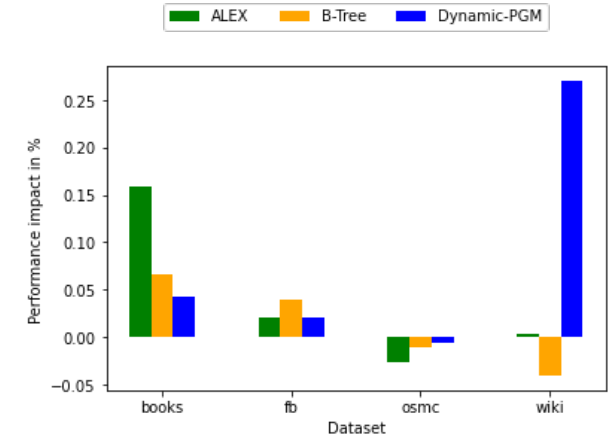
Non-poisoned scenario:



Poisoned scenario:

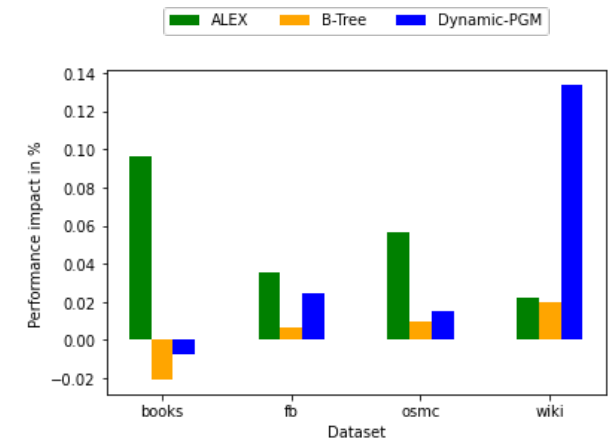




Performance impact:



4. Write-only workload

20M inserts



Note: All datasets in this experiment consist of 200M uint64 keys. The benchmark was evaluated on a e2-standard-8 VM instance running on Google Compute Cloud (8 vCPUs, 32GB RAM). For each index, multiple configurations were tested.

Agenda

- 1** Introduction
- 2** Background on Adversarial ML
- 3** Poisoning Attacks on Learned Index Structures
- 4** Experimental Setup
- 5** Experimental Results
- 6** Conclusion

Research has shown that “adversarial workload matters”: Up to 20% performance deterioration observed in learned indexes



Limitations

- Poisoning attack executed with poisoning threshold of $p = 0.0001$ due to computational constraints (vs. $p = 0.01$ to $p = 0.2$)
- Experiments performed on Virtual Machine instance (*e2-standard-8*), thus sensitive to I/O demand on physical host

Recommendations

- Various approaches for defending against poisoning attacks on linear regression (TRIM¹, LID²), but not directly applicable
- In LIS, the rank for each key depends on value of all other keys (mitigations needs to iteratively re-calibrate!)
- Removal of poisoning keys often difficult (concentrated around legitimate keys)

Future Directions

- Assess other models for constructing learned index structures, such as polynomial interpolation³ or logarithmic error regression⁴
- Leverage the update functionality of LIS to insert / delete keys at runtime

1) Jagielski, M., et al. (2018). *Manipulating machine learning: Poisoning attacks and countermeasures for regression learning*. 2018 IEEE Symposium on Security and Privacy (SP), IEEE

2) Weerasinghe, S., et al. (2020). *Defending regression learners against poisoning attacks*. arXiv preprint arXiv:2008.09279.

3) Setiawan, N. F., et al. (2020). *Function interpolation for learned index structures*. Australasian Database Conference, Springer.

4) Eppert, M., et al. (2021). *A Tailored Regression for Learned Indexes: Logarithmic Error Regression*. Fourth Workshop in Exploiting AI Techniques for Data Management.