IPDPS
St. Petersburg, Florida USA
May 15-19, 2023

CHARLOTTE
COLLEGE OF COMPUTING AND INFORMATICS

Drill: Log-based Anomaly Detection for Large-scale Storage Systems Using Source Code Analysis

Di Zhang[1], Chris Egersdoerfer[1], Tabassum Mahmud[2], Mai Zheng[2], and Dong Dai[1]
[1]University of North Carolina at Charlotte
[2]Iowa State University

DIRLAB

Outline

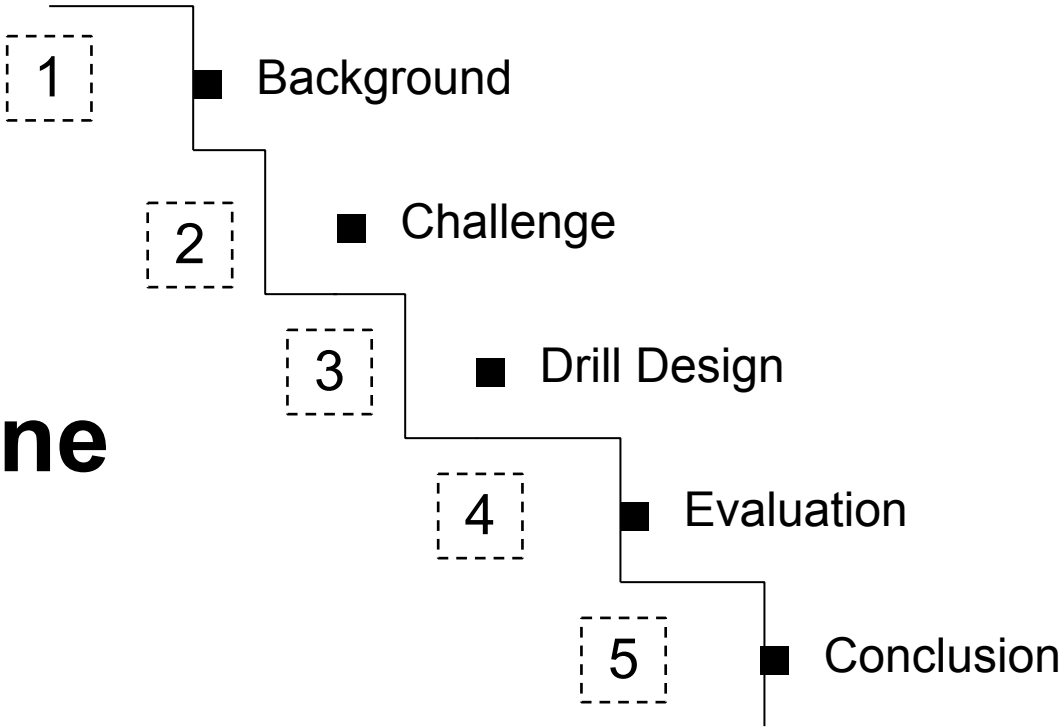
1
Background
2
Challenge
3
Drill Design
4
Evaluation
5
Conclusion

Outline

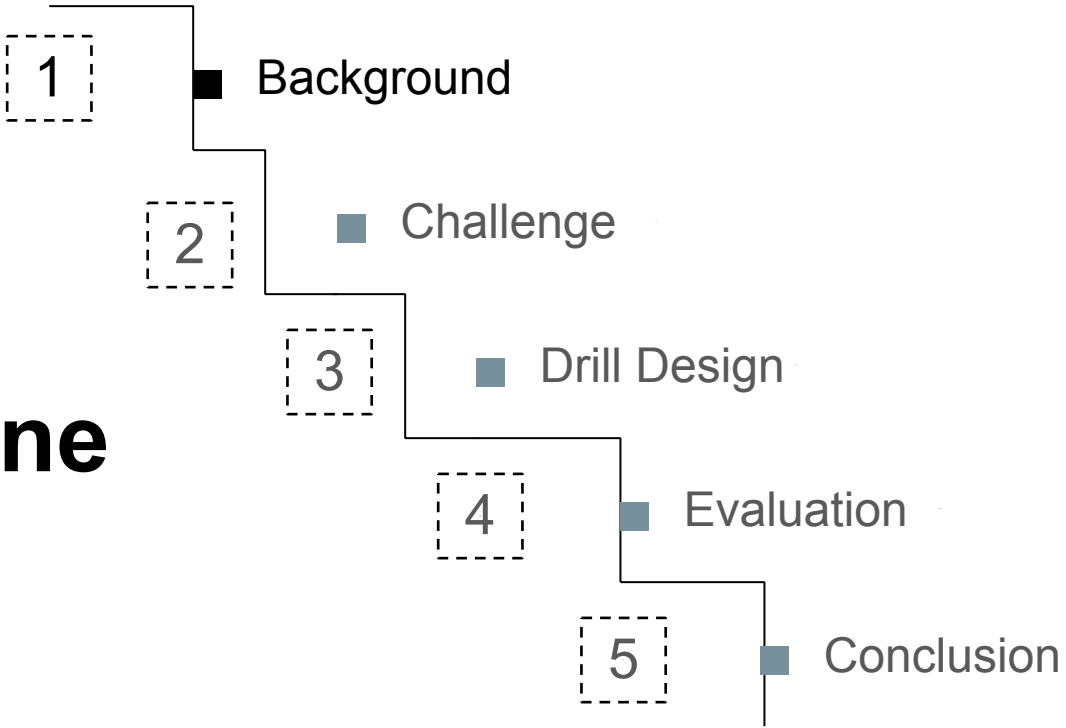
1
Background
2
Challenge
3
Drill Design
4
Evaluation
5
Conclusion

Importance of Anomaly Detection

Fugaku

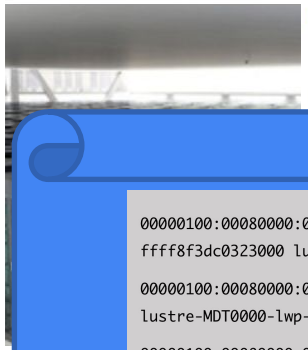

Summit

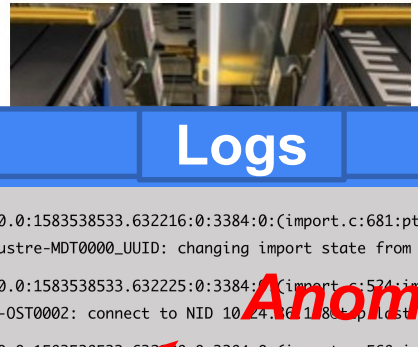

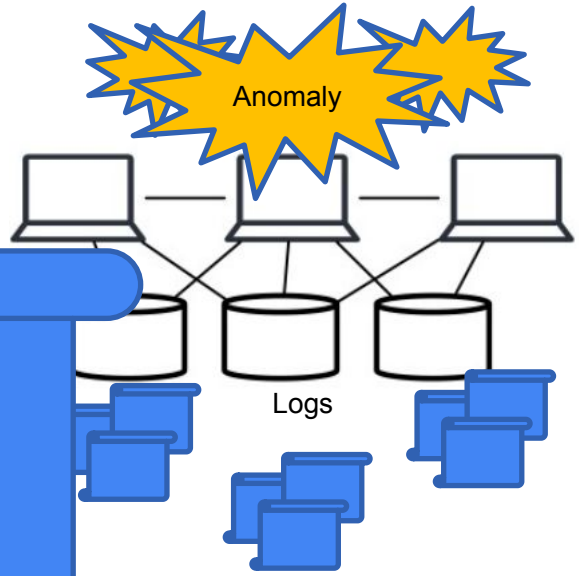

Logs

```
00000100:00080000:0.0:1583538533.632216:0:3384:0:(import.c:681:ptlrpc_connect_import())
ffff8f3dc0323000 lustre-MDT0000_UUID: changing import state from DISCONN to CONNECTING

00000100:00080000:0.0:1583538533.632225:0:3384:0:(import.c:524:import_select_connection())
lustre-MDT0000-lwp-OST0002: connect to NID 10.24.26.118@tcp last attempt 55-138959

00000100:00080000:0.0:1583538533.632228:0:3384:0:(import.c:568:import_select_connection())
lustre-MDT0000-lwp-OST0002: tried all connections, increasing latency to 50s

00000100:00080000:0.0:1583538558.632332:0:3384:0:(pinger.c:217:ptlrpc_pinger_process_import())
lustre-MDT0000-lwp-OST0002_UUID->lustre-MDT0000_UUID: level CONNECTING/4 force 0 force_next 0
deactive 0 pingable 1 suppress 0

00000100:00080000:0.0:1583538558.632342:0:3384:0:(pinger.c:230:ptlrpc_pinger_process_import())
lustre-MDT0000-lwp-OST0002_UUID->lustre-MDT0000_UUID: not pinging (in recovery or recovery
disabled CONNECTING)

00000100:00080000:0.0:1583538558.632228:0:3384:0:(import.c:568:import_select_connection())
lustre-MDT0000-lwp-OST0002: tried all connections, increasing latency to 60s
```

Anomaly!

BeeGFS®

ceph

DIRLAB

IPDPS
St. Petersburg, Florida USA
May 15-19, 2023

Workflow of Log-based Anomaly Detection

CHARLOTTE
COLLEGE OF COMPUTING AND INFORMATICS

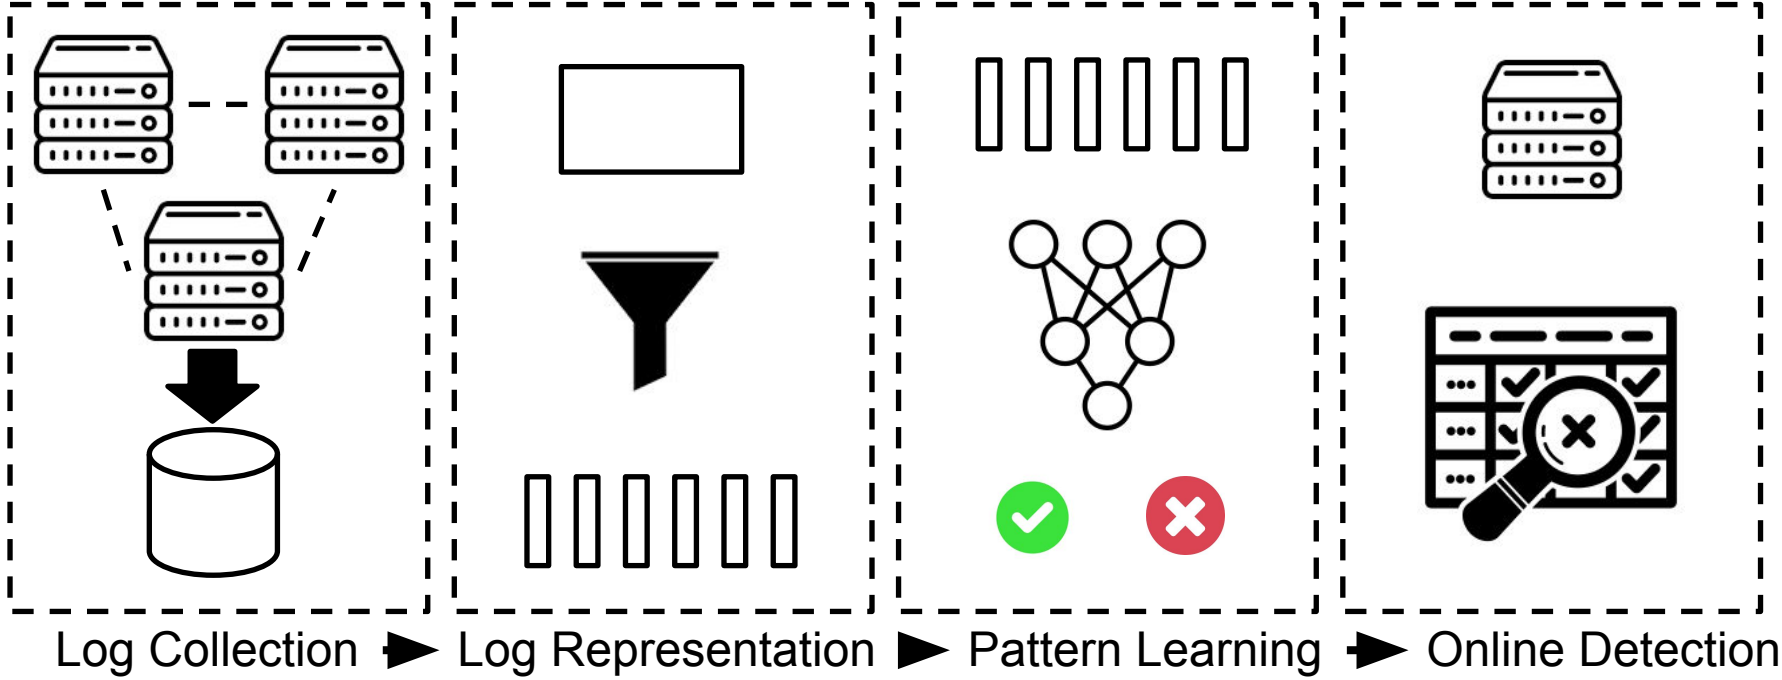

DIRLAB

IPDPS
St. Petersburg, Florida USA
May 15-19, 2023

Workflow of Log-based Anomaly Detection

CHARLOTTE
COLLEGE OF COMPUTING AND INFORMATICS

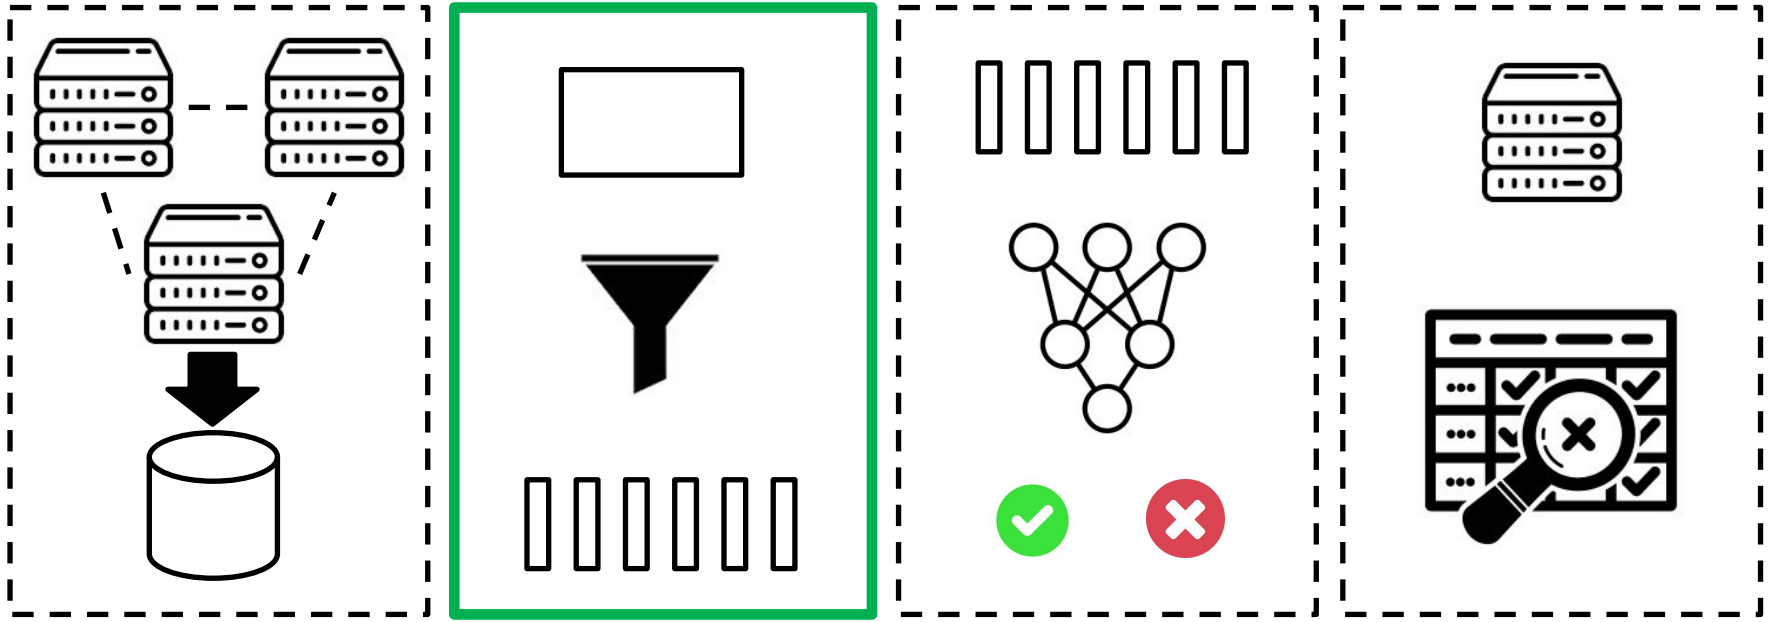

DIRLAB

Log Representation: Log Index

PCA
SOSP'09

Invariant Mining
ATC'10

DeepLog
CCS'17

00000100:00080000:0.0:1607448618.327577:0:2290:0:(recover.c:58:ptlrpc_initiate_recovery()) lustre-OST0000_UUID: starting recovery

00000100:00080000:0.0:1607448618.327580:0:2290:0:(import.c:681:ptlrpc_connect_import()) ffffa139cab87800 lustre-OST0000_UUID: changing import state from DISCONN to CONNECTING

00000100:00080000:0.0:1607448618.327589:0:2290:0:(import.c:524:import_select_connection()) lustre-OST0000-osc-MDT0000: connect to NID 10.0.0.8@tcp last attempt 4296114409

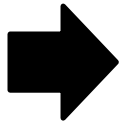

1

2

3

Log Representation: Log Content

LogAnomaly
IJCAI'19

NeuralLog
ASE'21

00000100:00080000:0.0:1607448618.327577:0:2290:0:(recover.c:58:ptlrpc_initiate_recovery()) lustre-OST0000_UUID: starting recovery

00000100:00080000:0.0:1607448618.327580:0:2290:0:(import.c:681:ptlrpc_connect_import()) ffffa139cab87800 lustre-OST0000_UUID: changing import state from DISCONN to CONNECTING

00000100:00080000:0.0:1607448618.327589:0:2290:0:(import.c:524:import_select_connection()) lustre-OST0000-osc-MDT0000: connect to NID 10.0.0.8@tcp last attempt 4296114409

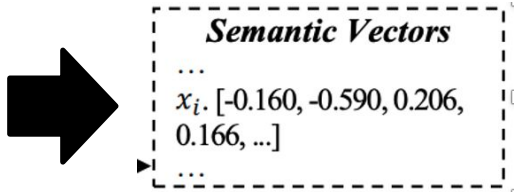

Outline

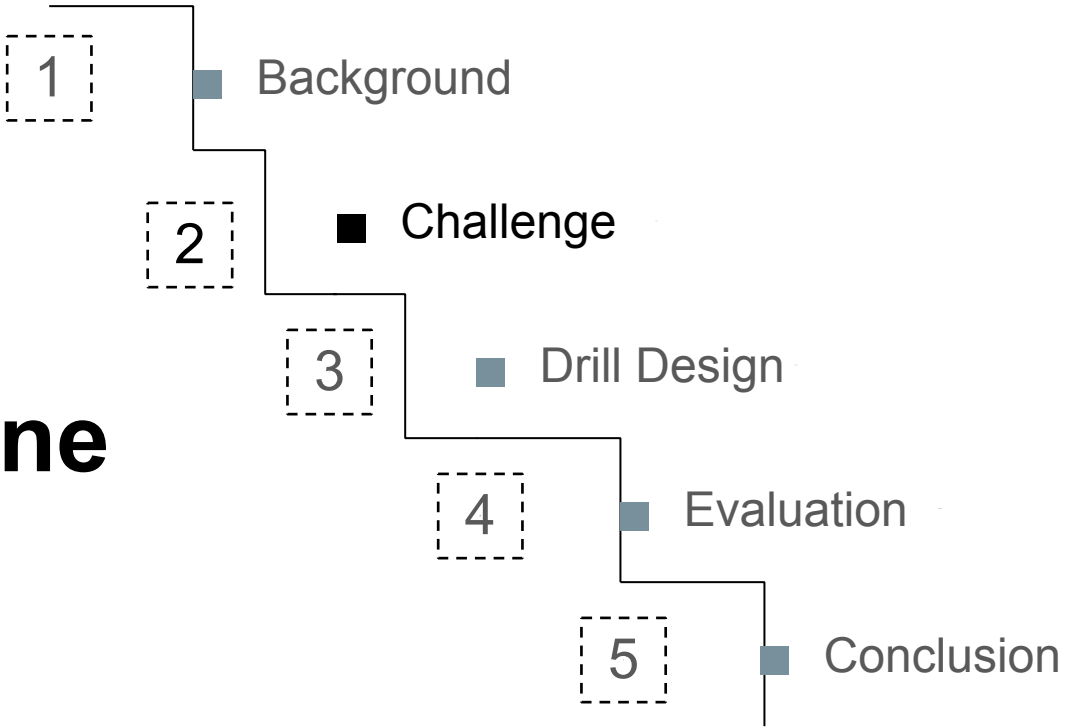
1
Background
2
Challenge
3
Drill Design
4
Evaluation
5
Conclusion

Challenge 1: Index-based Unseen Log Issues

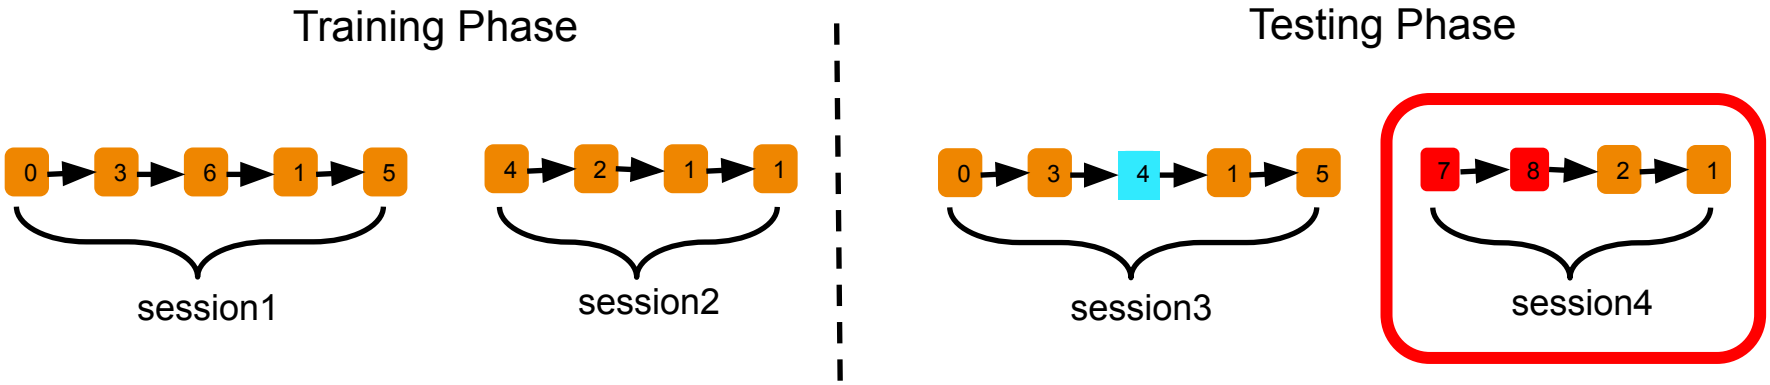

A dataset covers **29** unique runtime logs.
In total **956** distinct log statements in the source code.

Log index oversimplifies the representation of logs, missing important and valuable information.

Challenge 2: Log Content is Insufficient

INFO

1

00000100:00080000:0.0:1607448618.327577:0:2290:0:(recover.c:58:ptlrpc_initiate_recovery())
lustre-OST0002-osc-MDT0000: Wrote last used FID: [0x100020000:0x316f:0x0], index 2: 0

Log Index

Log Level

Log Content

Not Pure Natural Language

```
...
rc = kgnilnd_find_and_cancel_dgram(peer->gnp_net->gnn_dev,
        peer->gnp_nid);
if (rc) {
    LCONSOLE_INFO("Received NAK from %s for %s errno %d; "
        "canceled pending connect request\n",
        libcfs_nid2str(connreq->gncr_srcnid),
        libcfs_nid2str(connreq->gncr_dstnid), errno);
}
...
```

Log Context

The context of a log statement in the source code could be a strong feature in understanding the nature of the relevant logs

Outline

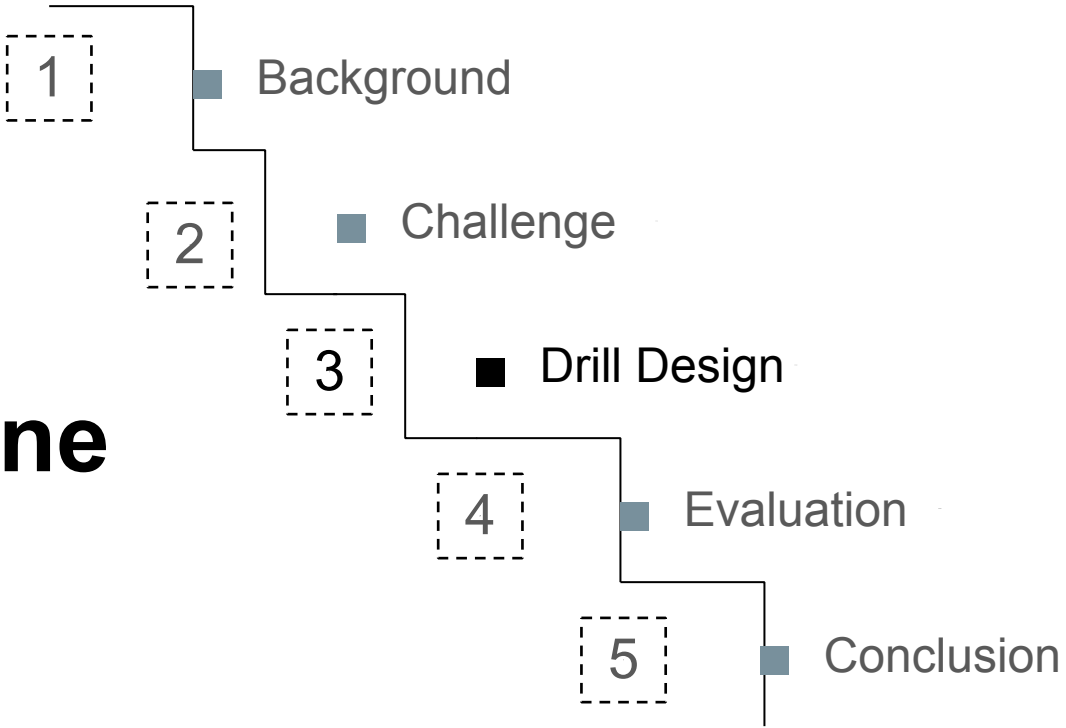
1
Background
2
Challenge
3
Drill Design
4
Evaluation
5
Conclusion

Drill Overview

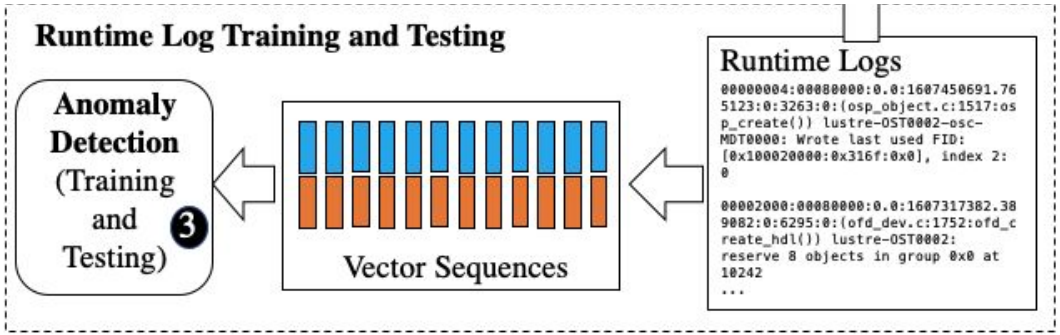

Sentiment Feature Builder: Motivation

```
...
CNETERR("Connection to %s at host %pI4h on port %d was "
        "refused: check that Lustre is running on that node.\n",
        libcfs_nid2str(peer_nid), &peer_ip, peer_port);
...
CDEBUG(D_HA, "recovery of %s on %s failed (%d)\n",
             obd2cli_tgt(imp->imp_obd),
             (char *)imp->imp_connection->c_remote_uuid.uuid, rc);
```

Negative

```
CDEBUG(D_HA, "%s: reserve %d objects in group %#llx"
             " at %llu\n", ofd_name(ofd),
             count, seq, next_id);
...
CDEBUG(D_HA, "%s: Wrote last used FID: "DFID", index %d: %d\n",
             d->opd_obd->obd_name, PFID(fid), d->opd_index, rc);
```

Neutral

```
CDEBUG(D_NET,"%s Route resolved: %d\n",
               libcfs_nid2str(peer_ni->ibp_nid), event->status);
...
CDEBUG(D_HA, "%s: transno %lld is committed\n",
             ccb->llcc_tgt->lut_obd->obd_name, ccb->llcc_transno);
```

Positive

Sentiment Feature Builder: Design

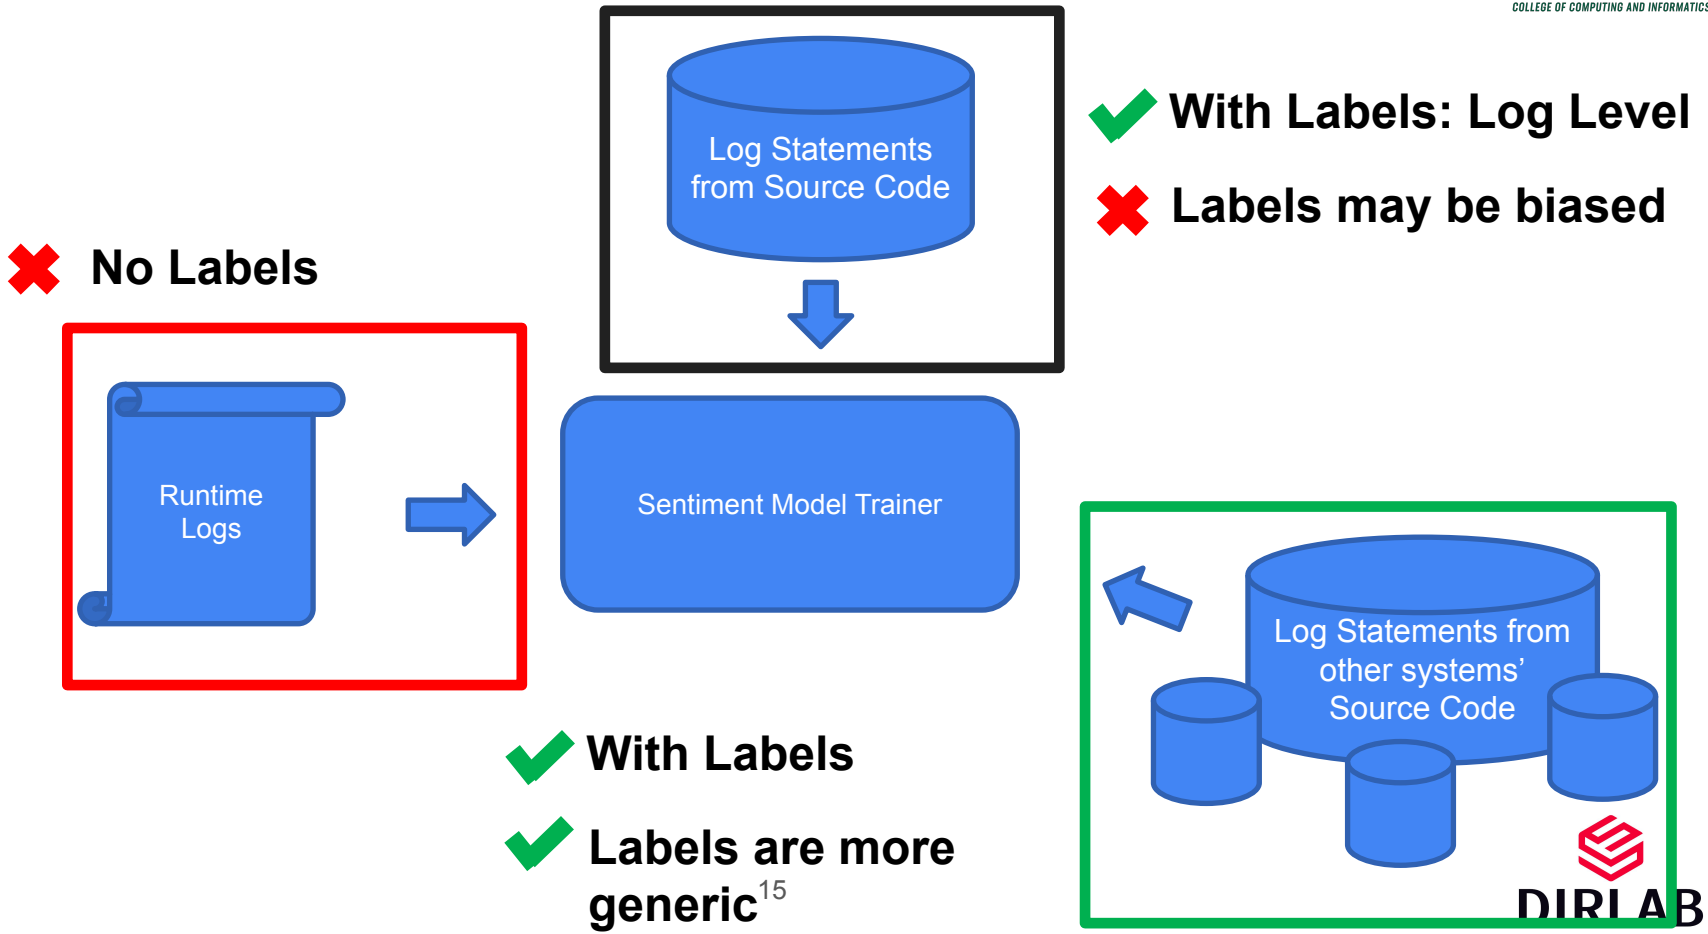

Context Feature Builder: Motivation

```
...
1  rc = ostid_set_id(&oa->o_oi, ostid_id(&oinfo->loi_oi));
2  if (rc) {
3      CERROR("Bad %llu to set " DOSTID " : rc %d\n",
4          (unsigned long long)ostid_id(&oinfo->loi_oi),
5          POSTID(&oa->o_oi), rc);
6  }
...
```

```
...
1  rc = kgnilnd_find_and_cancel_dgram(peer->gnp_net->gnn_dev,
2          peer->gnp_nid);
3  if (rc) {
4      LCONSOLE_INFO("Received NAK from %s for %s errno %d; "
5          "canceled pending connect request\n",
6          libcfs_nid2str(connreq->gncr_srcnid),
7          libcfs_nid2str(connreq->gncr_dstnid), errno);
8  }
...
```

```
rc = some_function();
if (rc) {
   log();
}
```

Context Feature Builder: Desi

```
...
1  rc = mdt_attr_get_complex(info, mo, ma);
2  if (rc) {
3      CERROR("file attribute read error for "DFID": %d.\n",
4          PFID(mdt_object_fid(mo)), rc);
5      RETURN(rc);
6  }
...
```

ReturnTypeI: *check rc*

MessageType: *variable checking*

ReturnTypeII : *immediately return*

ControlType: *if*

Algorithm 1: Context Feature Extraction

Input: IR, F
Output: Context Features

```
1  Function ExtractContextFeature(IR, F)
2      CreateCFG(IR, F)
3      foreach BB ∈ F do
4          foreach I ∈ BB do
5              if log_statement
                (CERROR || CDEBUG) then
6                  CheckControlType(BB)
7                  CheckMessageType(I)
8                  CheckReturnTypeI(BB)
9                  CheckReturnTypeII(I)
10 Function CheckControlType(BB)
11     // implemented by checking conditional jump
12     if is_conditional_Block(BB) then
13         control_type ← "if"
14     else
15         // detecting cycle using depth-first-search
16         if is_loop_Block(BB) then
17             control_type ← "loop"
18         else
19             control_type ← "null"
20 Function CheckMessageType(I)
21     if num_of_operands(I) > 1 then
22         message_type ← "variable_check"
23     else
24         message_type ← "status_monitor"
25 Function CheckReturnTypeI(BB)
26     if is_conditional_Block(BB) then
27         condition ← jump_condition
28         // back-trace the condition
29         if is_function_return_value(condition)
             then
30             log-after-return ← "yes"
31 Function CheckReturnTypeII(I)
32     if is_return_statement(next_instruction(I)) then
33         return-after-log ← "yes"
```

Outline

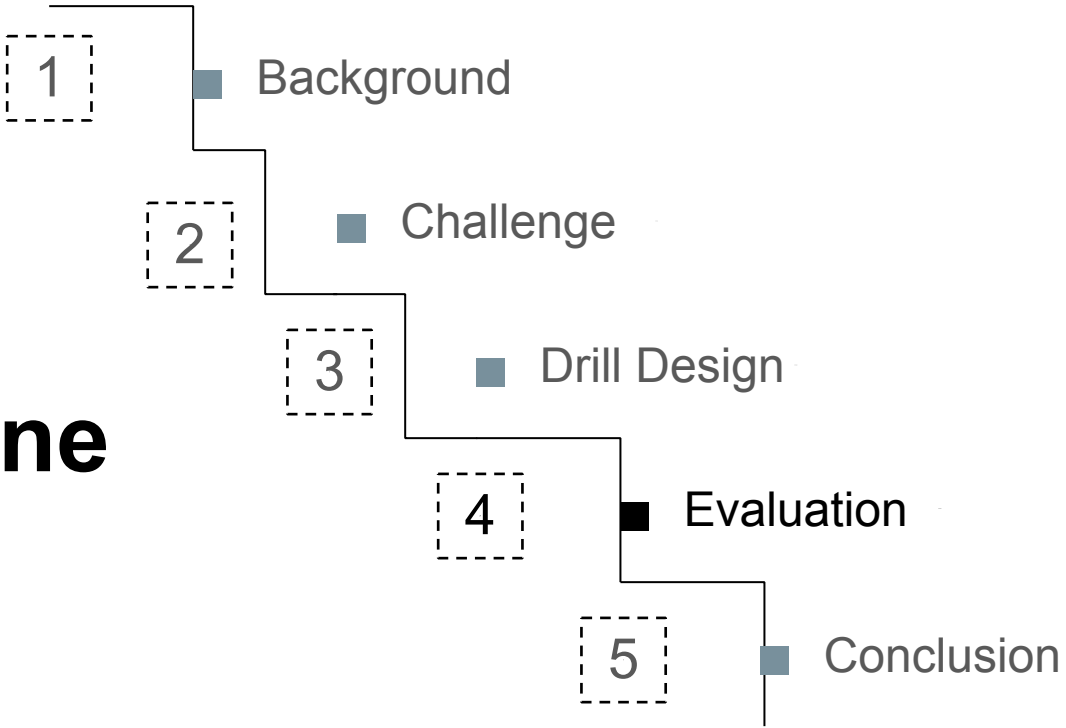
1
Background
2
Challenge
3
Drill Design
4
Evaluation
5
Conclusion

Evaluation Outline

- How is the impact of **sentiment feature?**
- How is the impact of **context feature?**

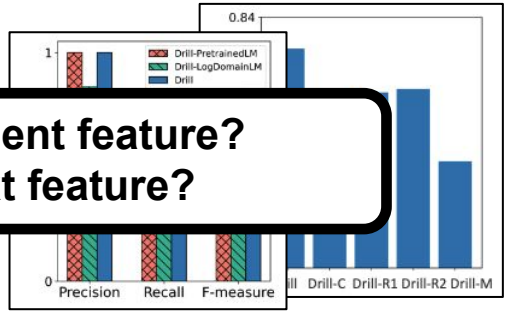

- How is the performance of Drill on **different systems**?
- How is the performance of Drill on **streaming analytics**?

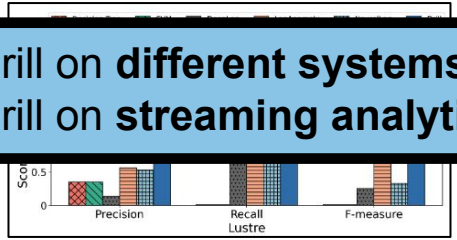

Is Drill robust on **partial training data**?

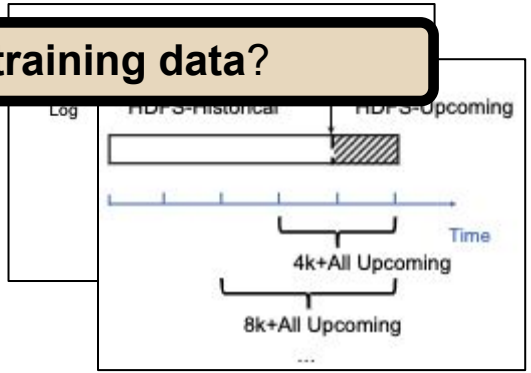

Evaluation Outline

- How is the impact of **sentiment feature?**
- How is the impact of **context feature?**

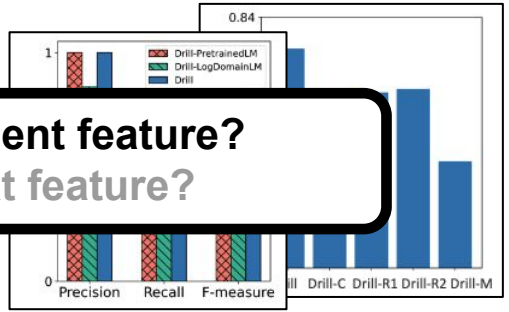

- How is the performance of Drill on **different systems**?
- How is the performance of Drill on **streaming analytics?**

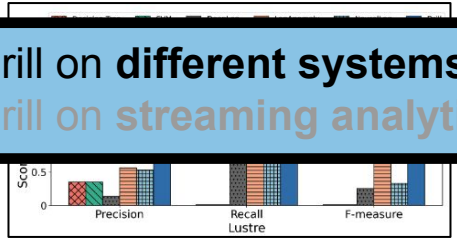

Is Drill robust on **partial training data?**

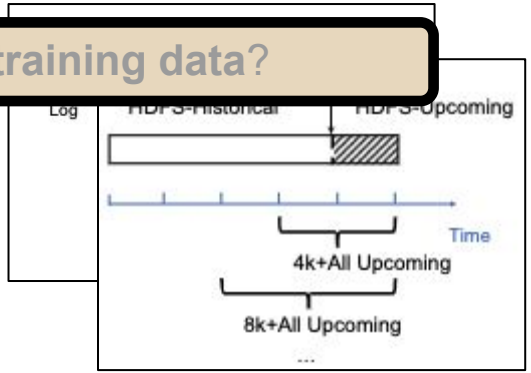

Impact of Sentiment Features

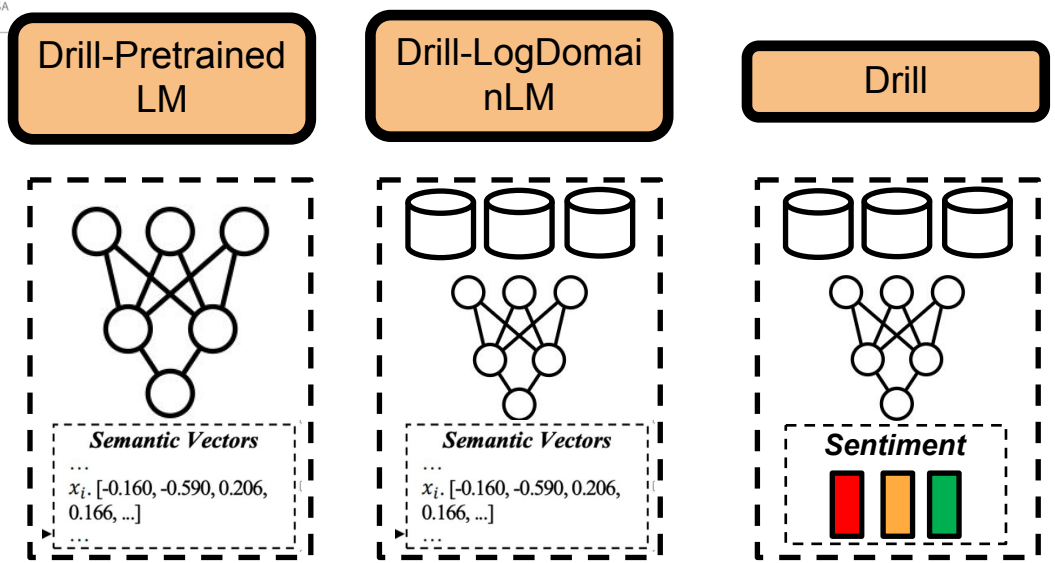

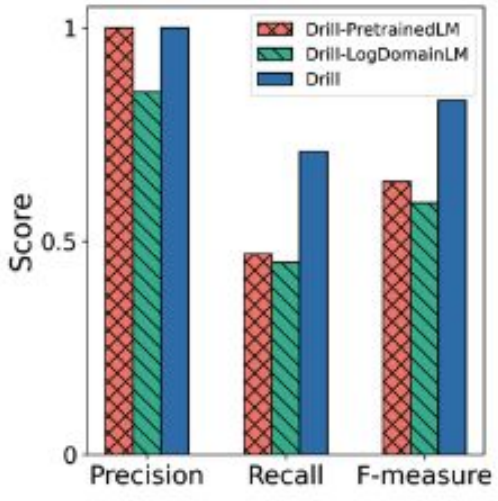

The volume of domain-specific data is not enough to finetune a workable language model

Drill on Different Systems

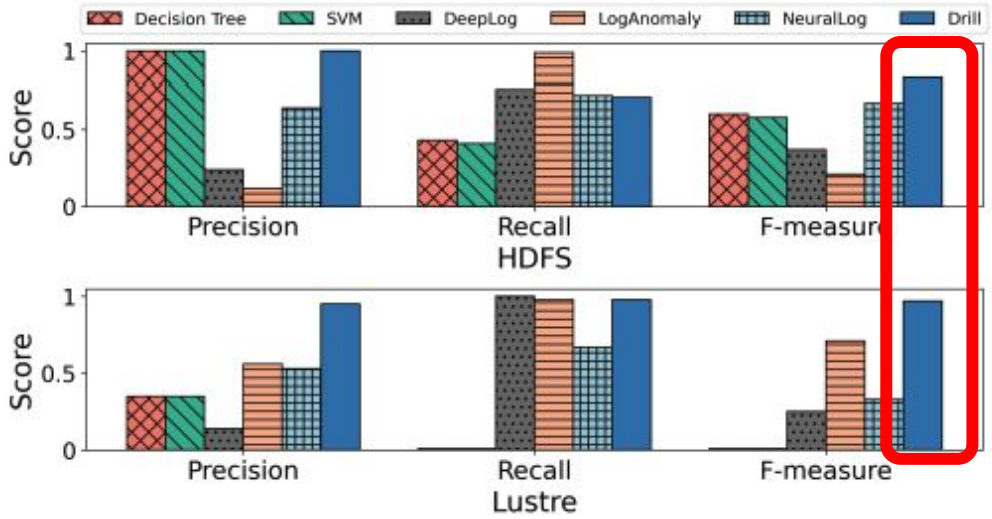

Drill achieves the best performance among all the six approaches on F-measure, presenting its effectiveness.

Outline

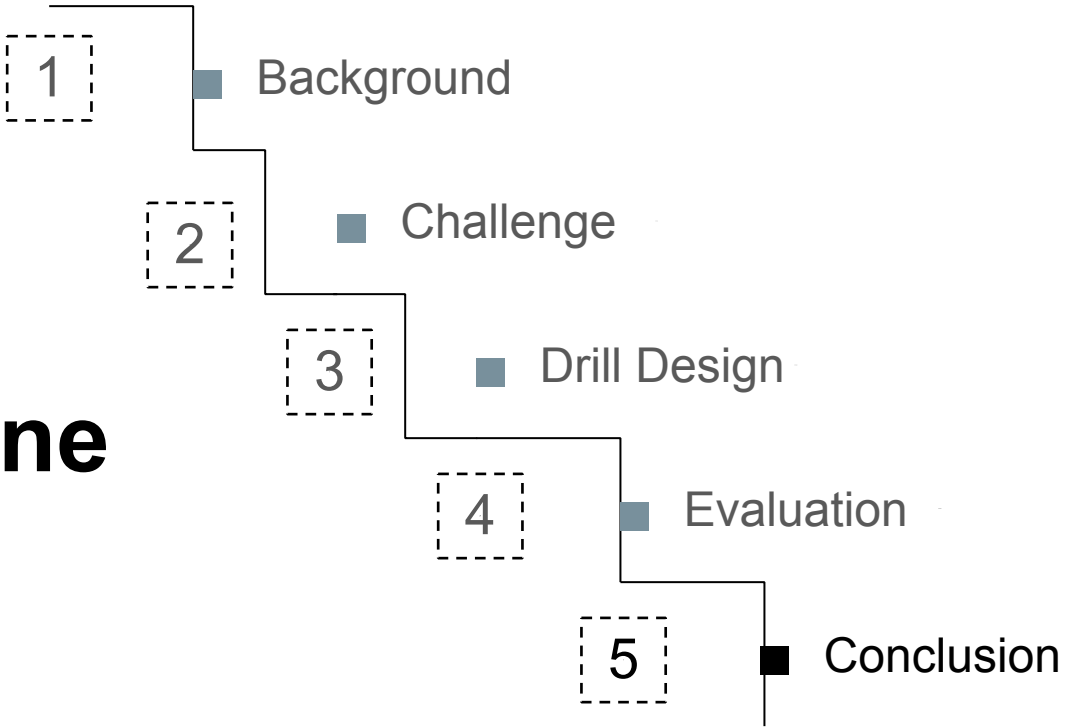
1
Background
2
Challenge
3
Drill Design
4
Evaluation
5
Conclusion

Conclusion and Future Work

- Conclusion:
  - We propose to use a storage system-specific sentiment language model and context-based feature extraction to detect the anomaly and show its effectiveness.
  - Our evaluations show Drill outperforms state-of-the-art approaches on two representative large-scale storage systems, HDFS and Lustre.

- Future Work:
  - Explore the possibility to consider more features besides the log statement description.
  - Apply more sophisticated language models, such as BERT for sentiment analysis.

Github Repo
[Code]
[Dataset]

Q&A
Thank you!

Backup: Use ChatGPT to detect anomaly in log

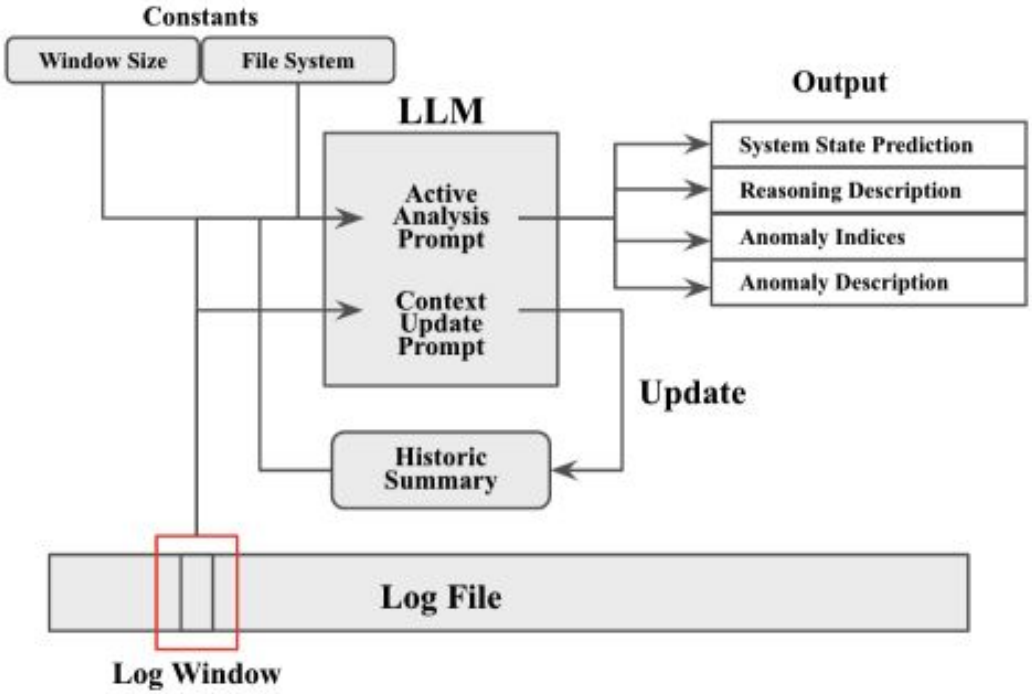

Backup: Use ChatGPT to detect anomaly in log

Log Lines

00000004:00080000:0.0:1607317389.467053:0:5483:0:(osp_object.c:1517:osp_create()) lustre-OST0000-osc-MDT0000: Wrote last used FID: [0x100000000:0x4c2b:0x0], index 0: 0 ||
00000004:00080000:0.0:1607317389.467055:0:5483:0:(osp_object.c:1517:osp_create()) lustre-OST0001-osc-MDT0000: Wrote last used FID: [0x100010000:0x4c2b:0x0], index 1: 0 ||
00000004:00080000:0.0:1607317389.467056:0:5483:0:(osp_object.c:1517:osp_create()) lustre-OST0002-osc-MDT0000: Wrote last used FID: [0x100020000:0x4c6b:0x0], index 2: 0 ||
00080000:00020000:0.0:1607317389.469131:0:5423:0:(osd_handler.c:1588:osd_trans_commit_cb()) transaction @0xffff9676bae890c0 commit error: 2 ||
00080000:00020000:0.0:1607317389.469136:0:5423:0:(osd_handler.c:1588:osd_trans_commit_cb()) transaction @0xffff9676c8e490c0 commit error: 2 ||
00080000:00020000:0.0:1607317389.469137:0:5423:0:(osd_handler.c:1588:osd_trans_commit_cb()) transaction @0xffff9676ba5393c0 commit error: 2 ||
00080000:00020000:0.0:1607317389.469138:0:5423:0:(osd_handler.c:1588:osd_trans_commit_cb()) transaction @0xffff9676c8e49900 commit error: 2 ||
00080000:00020000:0.0:1607317389.469138:0:5423:0:(osd_handler.c:1588:osd_trans_commit_cb()) transaction @0xffff9676c8e49540 commit error: 2 ||
00080000:00020000:0.0:1607317389.469139:0:5423:0:(osd_handler.c:1588:osd_trans_commit_cb()) transaction @0xffff9676c8e49b40 commit error: 2 ||
00080000:00020000:0.0:1607317389.469139:0:5423:0:(osd_handler.c:1588:osd_trans_commit_cb()) transaction @0xffff9676c8e49e40 commit error: 2 ||

LLM Anomaly Description

Historic Summary: Continuing from the previous logs, Lustre-OST0000, Lustre-OST0001, and Lustre-OST0002 continue to write their last used FIDs to MDT0000 with incremented indexes. Additionally, multiple transactions with commit errors are recorded in the logs. No other significant events are noted.
system health prediction: Critical
description of system health: The multiple transactions with commit errors indicate a problem with the Lustre file system's ability to commit changes. This could lead to data loss or corruption if not addressed immediately.
Anomalies: [0 0 0 1 1 1 1 1 1 1]
Description: The commit errors in the transactions indicate that Lustre is having issues writing changes to disk. This could cause data inconsistencies or loss and is a critical issue that needs to be addressed.

Backup: Use ChatGPT to detect anomaly in log

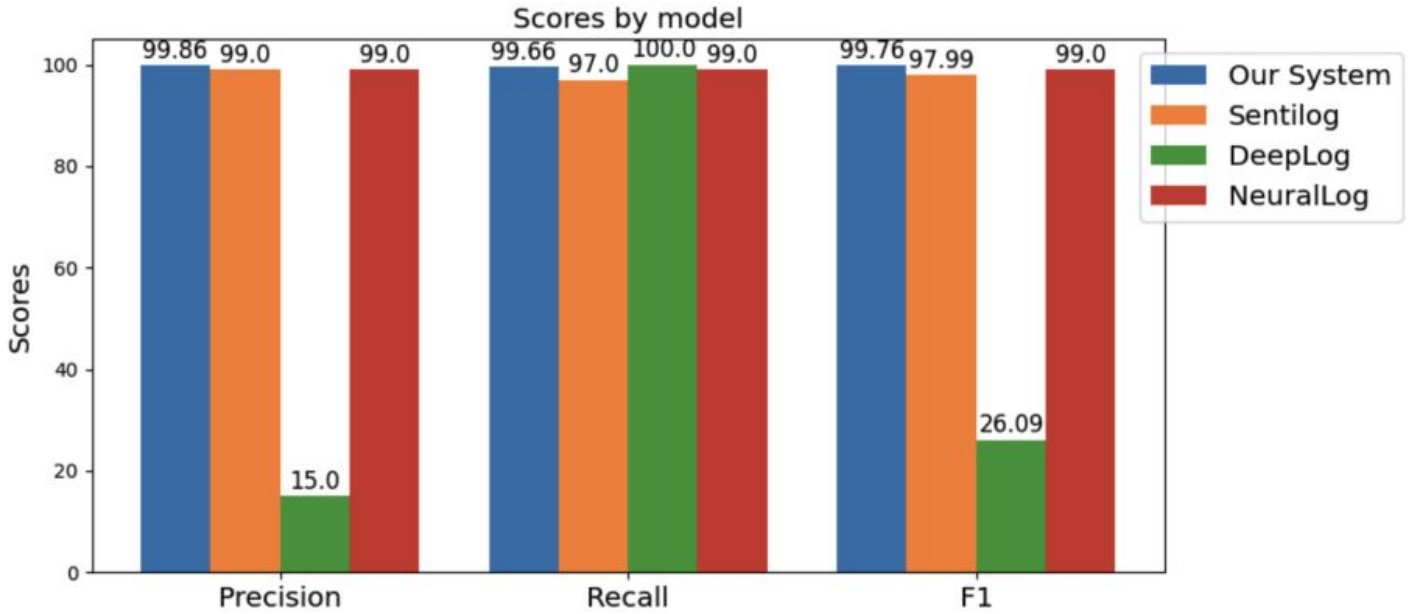

Existing Work: Two Different Ways

Log Index

Log Content

DeepLog

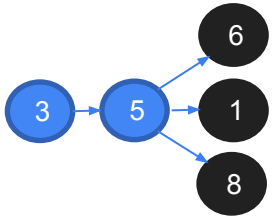

Search keyword,
e.g., “error”, “exception”.
Compare synonyms and antonyms,
e.g., LogAnomaly
BERT-based language model.
e.g., NeuralLog

IM

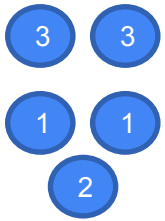

# of 1 = 2
# of 2 = 1
# of 3 = 2
Invariant:
# of 1 == # of 3

…

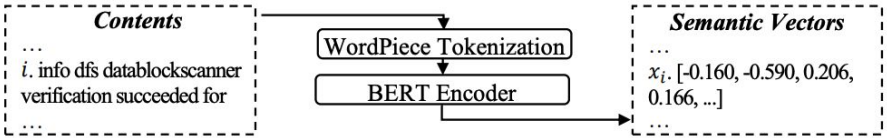

Sentiment Feature Builder: Design