

TOWARDS UNDERSTANDING AND ANALYZING RATIONALE IN COMMIT MESSAGES USING A KNOWLEDGE GRAPH APPROACH

Mouna Dhaouadi[†], **Bentley James Oakes^{‡†}**, Michalis Famelis[†]

[†] Université de Montréal
[‡] Polytechnique Montréal

MDEIntelligence Workshop
at MODELS 2023




**POLYTECHNIQUE
MONTRÉAL**

UNIVERSITÉ
D'INGÉNIERIE



Mouna Dhaouadi




Bentley James Oakes



Michalis Famelis



Mouna Dhaouadi



Bentley James Oakes



Michalis Famelis

PhD topic :

Rationale extraction and management

For a **decision**, software/system developer has **rationale**
'Why' reasoning for their decision

For a **decision**, software/system developer has **rationale**
'Why' reasoning for their decision

Rationale is *useful information* to deeply understand the system
Learn from mistakes, reuse solutions

Scope: Code commit messages

Scope: Code commit messages

‘‘Fix: Added check to prevent null pointer exception.’’

Scope: Code commit messages

‘‘Fix: Added check to prevent null pointer exception.’’
versus
‘‘changes’’

Scope: Code commit messages

‘‘Fix: Added check to prevent null pointer exception.’’
versus
‘‘changes’’

Challenges:

- What is rationale *exactly*?
- Implicit or unrecorded

Scope: Code commit messages

‘‘Fix: Added check to prevent null pointer exception.’’
versus
‘‘changes’’

Challenges:

- What is rationale *exactly*?
- Implicit or unrecorded
- Ambiguous language
 - Especially past/future tenses
 - Non-native writers
- Need technical understanding

Scope: Code commit messages

‘‘Fix: Added check to prevent null pointer exception.’’
versus
‘‘changes’’


Challenges:

- What is rationale *exactly*?
- Implicit or unrecorded
- Ambiguous language
 - Especially past/future tenses
 - Non-native writers
- Need technical understanding


Research question:

How can we *structure* and *extract* rationale information from code commits?

“Kantara: a framework for end-to-end rationale reconstruction”



“Kantara: a framework for end-to-end rationale reconstruction”



Purposes:

- As research to gain insight into rationale itself
- Allow developers to better understand presence of rationale in their software
 - Who, what, where, when
- Decide if a code commit has *insufficient rationale*
 - Flag or reject

1 Introduction

2 **Running Example**

3 Creating a Dataset

4 Automatic Sentence Classification

5 Structuring Rationale Information

6 Analyzing Rationale Information

7 Conclusion

Linux kernel:

- Development is through Git commits
- Culture for motivating/describing changes
- Valuable (highest-quality?) source of rationale

Linux kernel:

- Development is through Git commits
- Culture for motivating/describing changes
- Valuable (highest-quality?) source of rationale

Out-of-Memory Killer subsystem:

- When Linux kernel runs out of memory
- OOM-Killer is called to avoid crashing
- Two broad steps:
 - 1 Select “best” task to kill
 - 2 Force task to release memory and exit
- Meaningful heuristics

COMMIT EXAMPLE

#	Sentence
0	signal: Use SEND_SIG_PRIV not SEND_SIG_FORCED with SIGKILL and SIGSTOP

#	Sentence
0	signal: Use SEND_SIG_PRIV not SEND_SIG_FORCED with SIGKILL and SIGSTOP
1	Now that siginfo is never allocated for SIGKILL and SIGSTOP there is no difference between SEND_SIG_PRIV and SEND_SIG_FORCED for SIGKILL and SIGSTOP.
2	This makes SEND_SIG_FORCED unnecessary and redundant in the presence of SIGKILL and SIGSTOP.
3	Therefore change users of SEND_SIG_FORCED that are sending SIGKILL or SIGSTOP to use SEND_SIG_PRIV instead.
4	This removes the last users of SEND_SIG_FORCED.

- 1 Introduction
- 2 Running Example
- 3 **Creating a Dataset**
- 4 Automatic Sentence Classification
- 5 Structuring Rationale Information
- 6 Analyzing Rationale Information
- 7 Conclusion

Goal: Build a dataset of OOM-Killer commit sentences,
classified wrt rationale

Goal: Build a dataset of OOM-Killer commit sentences,
classified wrt rationale

Motivation:

- Gain insights through classification
- Empirical analysis
- Ground truth for automated classification

Goal: Build a dataset of OOM-Killer commit sentences,
classified wrt rationale

Motivation:

- Gain insights through classification
- Empirical analysis
- Ground truth for automated classification

Procedure:

- Collect 410 commits (paper reports 180)
- Remove merge commits, filter code sentences
- Three authors label sentences
- Resolve conflicts in discussion

Label	Meaning
Inapplicable	Pre-processing error or bad sentences (i.e., does not contain English sentences)
Supporting Facts	A narration of facts used to support a decision
Rationale	Reason for a decision or value judgment
Decision	An action or a change that has been made, including a description of the patch behaviour

Label	Meaning
Inapplicable	Pre-processing error or bad sentences (i.e., does not contain English sentences)
Supporting Facts	A narration of facts used to support a decision
Rationale	Reason for a decision or value judgment
Decision	An action or a change that has been made, including a description of the patch behaviour

- Obtained through piloting rounds and discussions
- **Multiple classifications per sentence**
- In disagreement, take classification union
- Fleiss kappa: Around 0.65 (fair to good agreement)

CLASSIFICATION EXAMPLE

- 0 Decision signal: Use SEND_SIG_PRIV not SEND_SIG_FORCED with SIGKILL and SIGSTOP

CLASSIFICATION EXAMPLE

- 0 **Decision** signal: Use SEND_SIG_PRIV not SEND_SIG_FORCED with SIGKILL and SIGSTOP
- 1 **Supporting Fact & Rationale** Now that siginfo is never allocated for SIGKILL and SIGSTOP there is no difference between SEND_SIG_PRIV and SEND_SIG_FORCED for SIGKILL and SIGSTOP.
- 2 **Rationale** This makes SEND_SIG_FORCED unnecessary and redundant in the presence of SIGKILL and SIGSTOP.
- 3 **Decision** Therefore change users of SEND_SIG_FORCED that are sending SIGKILL or SIGSTOP to use SEND_SIG_PRIV instead.
- 4 **Decision** This removes the last users of SEND_SIG_FORCED.

Common structure:

Decision summary phase ,

Supporting facts , Rationale , Decisions

Common structure:

Decision summary phase ,

Supporting facts , Rationale , Decisions

Amount of rationale:

- **97.5%** of commits contain rationale
- About **40-50%** of sentences per commit contain rationale

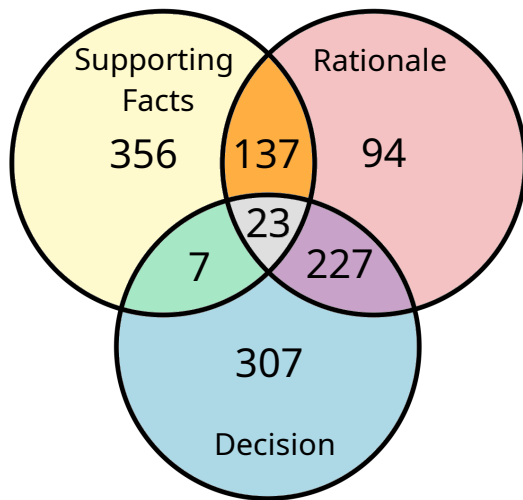
Common structure:

Decision summary phase ,

Supporting facts , Rationale , Decisions

Amount of rationale:

- **97.5%** of commits contain rationale
- About **40-50%** of sentences per commit contain rationale



- 1 Introduction
- 2 Running Example
- 3 Creating a Dataset
- 4 **Automatic Sentence Classification**
- 5 Structuring Rationale Information
- 6 Analyzing Rationale Information
- 7 Conclusion

AUTOMATIC CLASSIFICATION OF SENTENCES

Binary classification: Logistic regression, decision tree, SVM

Multi-label classification: Random Forest, XGBoost, KNN

Binary classification: Logistic regression, decision tree, SVM

Multi-label classification: Random Forest, XGBoost, KNN

XGBoost classification evaluation			
Label	Precision	Recall	F1-score
Decision	0.76	0.69	0.72
Rationale	0.62	0.41	0.49
Supporting Facts	0.64	0.68	0.66

Binary classification: Logistic regression, decision tree, SVM

Multi-label classification: Random Forest, XGBoost, KNN

XGBoost classification evaluation			
Label	Precision	Recall	F1-score
Decision	0.76	0.69	0.72
Rationale	0.62	0.41	0.49
Supporting Facts	0.64	0.68	0.66

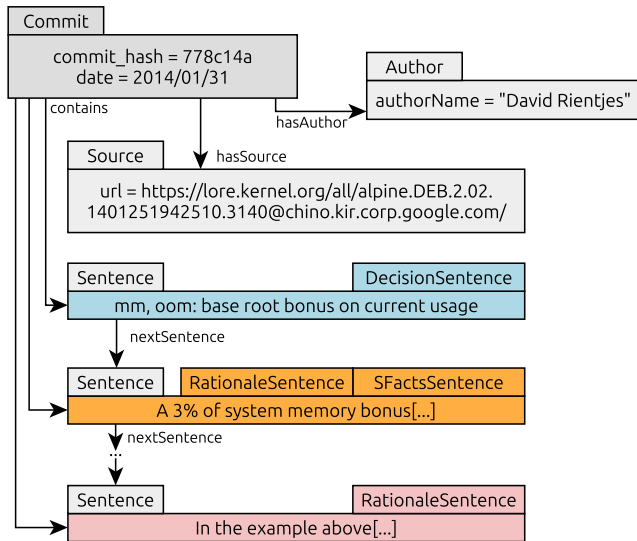
Insights:

- Overall poor performance
- Decisions easier to classify, rationale is harder

- 1 Introduction
- 2 Running Example
- 3 Creating a Dataset
- 4 Automatic Sentence Classification
- 5 Structuring Rationale Information**
- 6 Analyzing Rationale Information
- 7 Conclusion

RATIONALE INFO AS A GRAPH

- Prior work modelled relationships between commits
- Here, we model the *sentences* as a *knowledge graph*



Elaasar *et al.*, “openCAESAR: Balancing agility and rigor in model- based systems engineering,” in Proceedings of SAM Conference, 2023

Elaasar *et al.*, “openCAESAR: Balancing agility and rigor in model- based systems engineering,” in Proceedings of SAM Conference, 2023

Ontological Modelling Language (OML) for easily creating OWL ontologies and the knowledge graph

Elaasar *et al.*, "openCAESAR: Balancing agility and rigor in model- based systems engineering," in Proceedings of SAM Conference, 2023

Ontological Modelling Language (OML) for easily creating OWL ontologies and the knowledge graph

The screenshot displays the Rosetta IDE interface. The top menu bar includes File, Edit, Navigate, Search, Project, Run, Window, and Help. The left pane, titled 'Model Explorer', shows a tree view of the project structure: kepler16b-example (main), rationale_onto (Project Dependencies), build, gradle, src, and oml (geodes.iro.umontreal.ca). Under 'rationale', there are 'desc' and 'vocab' sub-projects. The 'vocab' sub-project contains 'authorV.owl', 'base.owl', 'rationaleVocab.owl', and 'vocabBundle.owl'. The right pane shows the OML code for 'rationaleVocab.owl' with the following content:

```

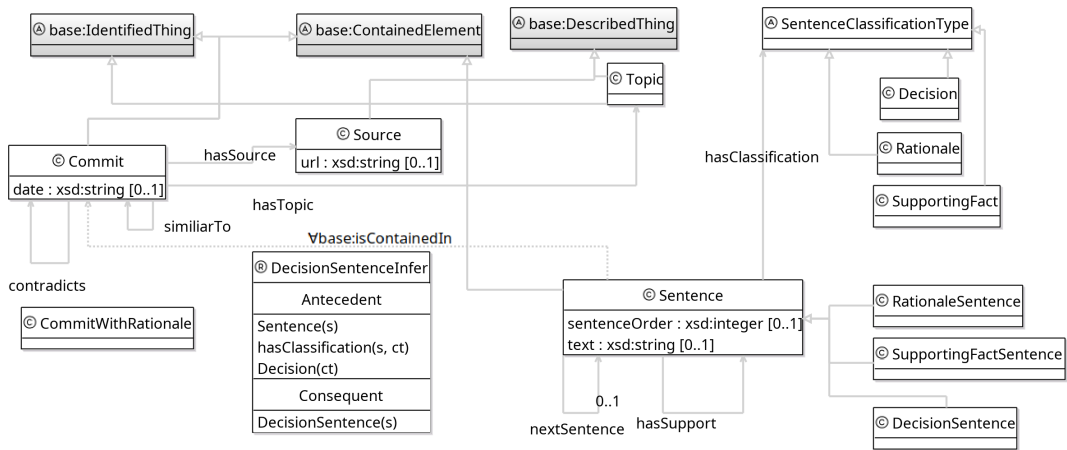
extends <https://geodes.iro.umontreal.ca/rationale/vocab/authorV#> as authorV

@rdfs:label "Commit"
@rdfs:comment "The core Commit concept"
concept Commit < base:ContainedElement, base:IdentifiedThing


@rdfs:label "Sentence"
@rdfs:comment "Sentence concept"
concept Sentence < base:ContainedElement [
    restricts all relation base:isContainedIn to Commit
]

scalar property text [
    domain Sentence
    range xsd:string
    functional
]
  
```

ONTOLOGY AND INFERRING



ONTOLOGY AND INFERRING



- Inferred multi-classification simplified analysis
- Easier to query for **CommitWithRationale**, **RationaleSentence**

- 1 Introduction
- 2 Running Example
- 3 Creating a Dataset
- 4 Automatic Sentence Classification
- 5 Structuring Rationale Information
- 6 Analyzing Rationale Information**
- 7 Conclusion

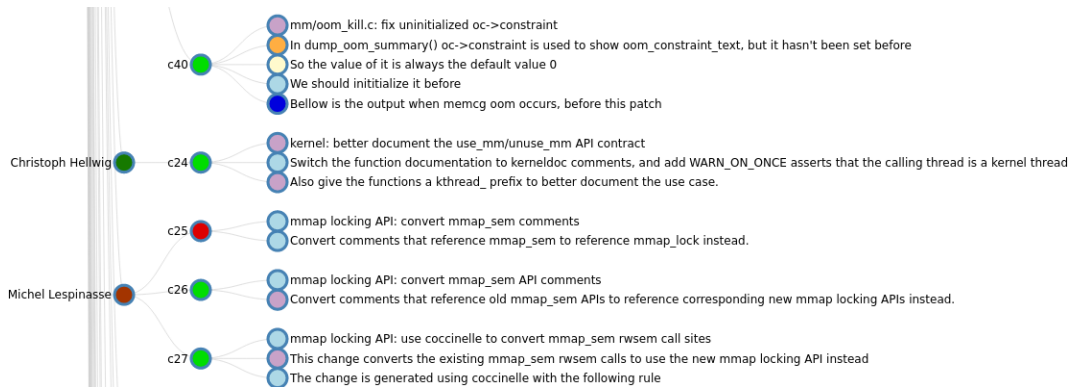
SPARQL queries on knowledge graph

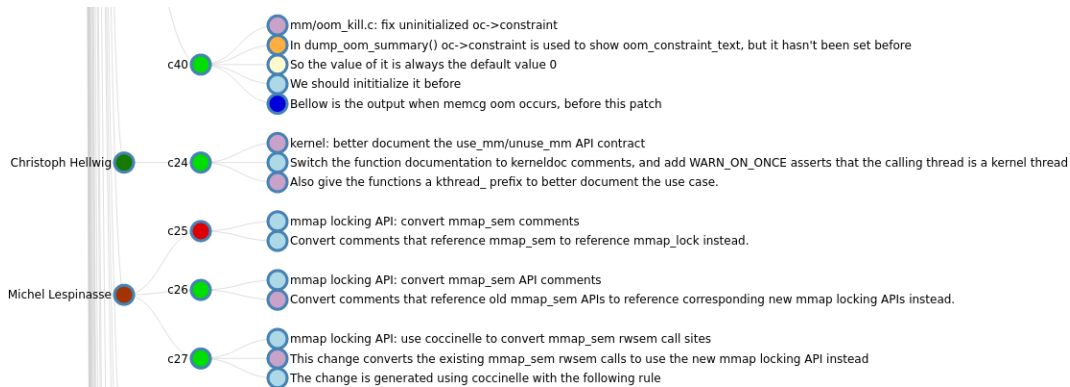
SPARQL queries on knowledge graph

Implemented:

- Listing authors and their commits
- List of sentences containing rationale
- List of commits containing rationale (uses inference)

```
{  
"author": { "value": "Michal Hocko" } ,  
"commit_id": { "value": "c0" } ,  
"text": { "value": "mm: reduce noise in show_mem for lowmem allocations" } ,  
"isCommitWithRationale": { "value": "true" } ,  
"isSentenceRationale": { "value": "true" } ,  
"isSentenceDecision": { "value": "true" } ,  
"isSentenceSupporting": { "value": "false" }  
}
```



Intention:

- Allow developers/researchers to understand rationale presence
- Identify commits/subsystems/developers without sufficient rationale

- 1 Introduction
- 2 Running Example
- 3 Creating a Dataset
- 4 Automatic Sentence Classification
- 5 Structuring Rationale Information
- 6 Analyzing Rationale Information
- 7 **Conclusion**

Summary:

- 1 Classification of sentences
 - Supporting fact, rationale, decision
- 2 Insights into rationale presence
 - 40% of commit, category overlap
- 3 Poor automatic classification
- 4 Structuring as knowledge graph, inferencing, and analysis
- 5 Visualization of rationale presence

Summary:

- 1 Classification of sentences
 - Supporting fact, rationale, decision
- 2 Insights into rationale presence
 - 40% of commit, category overlap
- 3 Poor automatic classification
- 4 Structuring as knowledge graph, inferencing, and analysis
- 5 Visualization of rationale presence

Future work:

- Finish dataset
- Investigate indicators for rationale
 - Evolution over time
 - Developer characteristics
- Improve auto. classification
- Kantara as a pull request bot

Summary:

- 1 Classification of sentences
 - Supporting fact, rationale, decision
- 2 Insights into rationale presence
 - 40% of commit, category overlap
- 3 Poor automatic classification
- 4 Structuring as knowledge graph, inferencing, and analysis
- 5 Visualization of rationale presence

Future work:

- Finish dataset
- Investigate indicators for rationale
 - Evolution over time
 - Developer characteristics
- Improve auto. classification
- Kantara as a pull request bot

Towards Understanding and Analyzing Rationale in Commit Messages using a Knowledge Graph Approach

Mouna Dhaouadi, Bentley James Oakes, Michalis Famelis

