Gradient-Free Optimization of Retrieval Augmented Prompts for Legality Classification of Clauses in German Employment Contracts

David Pauschert

October 07, 2024, Master’s Thesis Kick-off Presentation

Chair of Software Engineering for Business Information Systems (sebis)
Department of Computer Science
School of Computation, Information and Technology (CIT)
Technical University of Munich (TUM)
wwwmatthes.in.tum.de

Outline

Background & Motivation

Research Questions

Methodology

Timeline

© sebis

The master's thesis aims to optimize the AI-Assisted Legal Analysis and Correction of German Employment Contracts

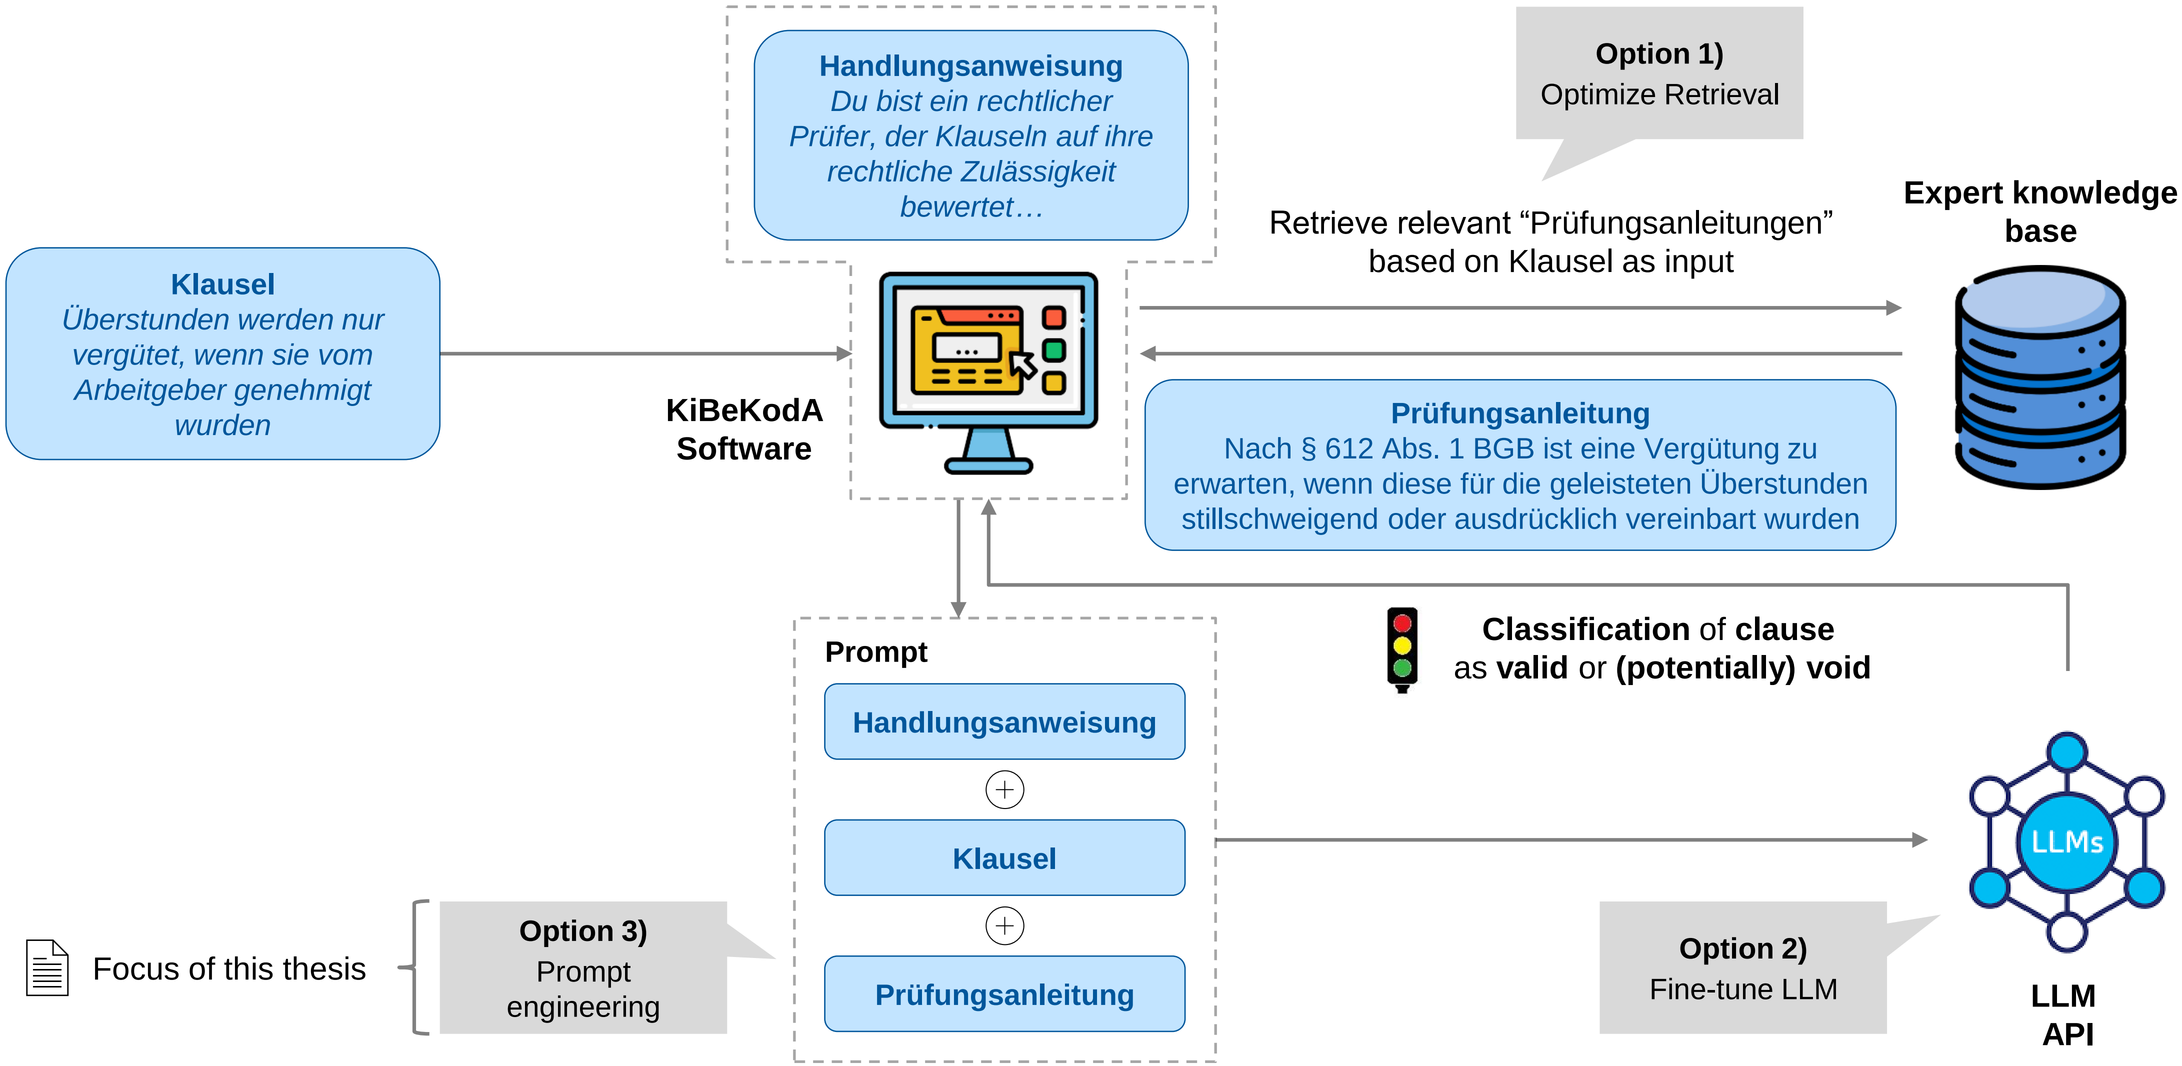

© sebis

Prompt optimization as the process of efficiently exploring the prompt space can be be performed both manually and (semi-)automatically

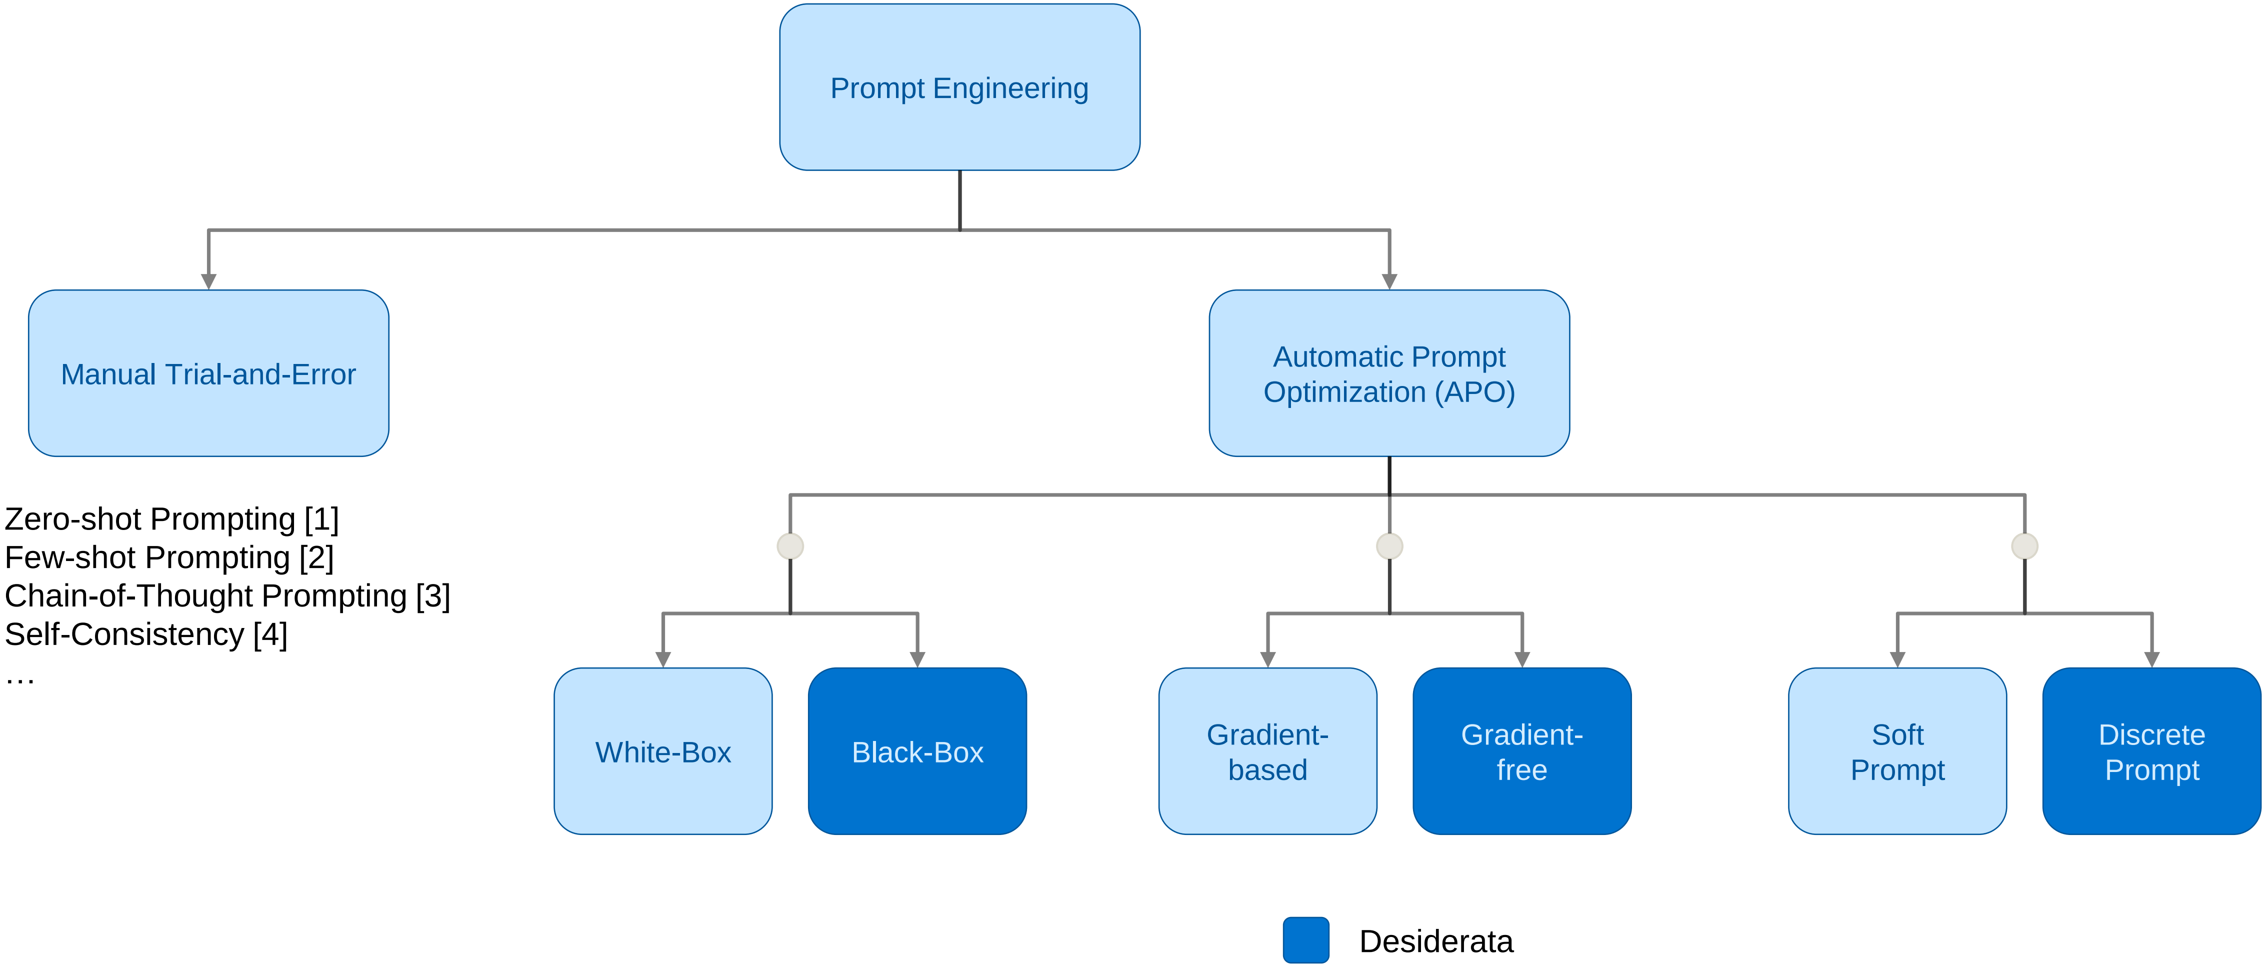

- Zero-shot Prompting [1]
- Few-shot Prompting [2]
- Chain-of-Thought Prompting [3]
- Self-Consistency [4]
- …

© sebis

Automatic Prompt Optimization (APO) Paradigm

$$P^*(x) = \arg\max_{P(\cdot)\sim\mathcal{P}} \mathbb{E}_{(x,y)\sim\mathcal{D}_{\text{val}}}\left[g\left(f_{\text{LLM}}(P(x)), y\right)\right]$$

Clause

Prompt Search Space

Score Function

True label

$$P = I \times E^k \times C^k$$

Instruction

Exemplars

Context (Prüfanleitung)

Instruction

Du bist ein deutscher Anwalt, welcher Klauseln aus Arbeitsverträgen unter Zuhiilfenahme von Prüfanleitungen hinsichtlich ihrer rechtlichen Zulässigkeit und Fairness bewertet

Clause

Ist die folgende Klausel unzulässig, unfair oder unbedenklich?
Überstunden werden nur vergütet, wenn sie vom Arbeitgeber genehmigt wurden

Context

Verwende folgende Prüfanleitung bei der Bewertung:
Nach § 612 Abs. 1 BGB ist eine Vergütung zu erwarten, wenn diese für die geleisteten Überstunden stillschweigend oder ausdrücklich vereinbart wurden

The challenge in engineering the prompt for the retrieved expert knowledge is that the content itself should not be changed

Ein Arbeitnehmer ist verpflichtet, ein ärztliches Attest vorzulegen, wenn die Krankheit länger als drei Kalendertage dauert, es sei denn, der Arbeitgeber verlangt das Attest bereits früher.

Ein Arbeitnehmer muss kein ärztliches Attest vorlegen, es sei denn, die Krankheit dauert länger als drei Kalendertage oder der Arbeitgeber fordert das Attest früher.

Die Pflicht zur Vorlage eines ärztlichen Attests bei Krankheit ist in § 5 Abs. 1 Entgeltfortzahlungsgesetz (EFZG) geregelt.

Which one will yield the best performance

© sebis

Outline

The master's thesis aims to investigate automatic optimization of expert-knowledge enhanced prompts for employment contract clause classification

	1 RQ 1	2 RQ 2	3 RQ 3
Research question	What is the **current state** of **automatic gradient-free prompt optimization** techniques and do they account for some part of the prompt containing **expert knowledge** where the **semantic** should **not** be **changed**?	How can these **methods** be **adapted** to optimize **prompts**, **augmented** with retrieved **legal expert knowledge**, for **classifying** the legality of German employment **contracts** without altering the semantics of the retrieved content?	How do the **adapted methods perform** on **legal clause classification** using clauses, labels and legal examination guides from the **Sylvenstein dataset** against other APO methods as baselines?
Methodology	Literature Review	Method development	Bench-Marking
Thesis Section	Background & Related Work	Proposed Method	Experimental Setup & Results

© sebis

Outline

Background & Motivation

Research Questions

Methodology & Initial Findings

Timeline

© sebis

RQ 1: Methodology & Initial Results

RQ 1 – Methodology & Initial Results

Methodology

Find relevant **APO method papers**[1] using backward and forward **snowballing** with [12] as starting paper for two iterations

Extract relevant **information** from **papers** using the following **research themes**:

1) Bibliographic Information
2) Optimization Scope (Instruction, examples, …)
3) Optimization Approach (Initialization, Evolution, Update, Termination)
4) Evaluation & Benchmark (task, data, model, metric)

Find meaningful **thematic clusters** for APO methods in scope

APO paradigms

	Characteristics		
APO Paradigms	Black-Box	Gradient-Free	Discrete
Soft Prompt Tuning	✗	✗	✗
Reinforcement Learning	✗	✗	✓
Bayesian Optimization	✓	✗	✓
LLM-based Optimizers	✓	✓	✓

Focus

The anatomy on an LLM based APO algorithm

RQ 1 – Initial Results

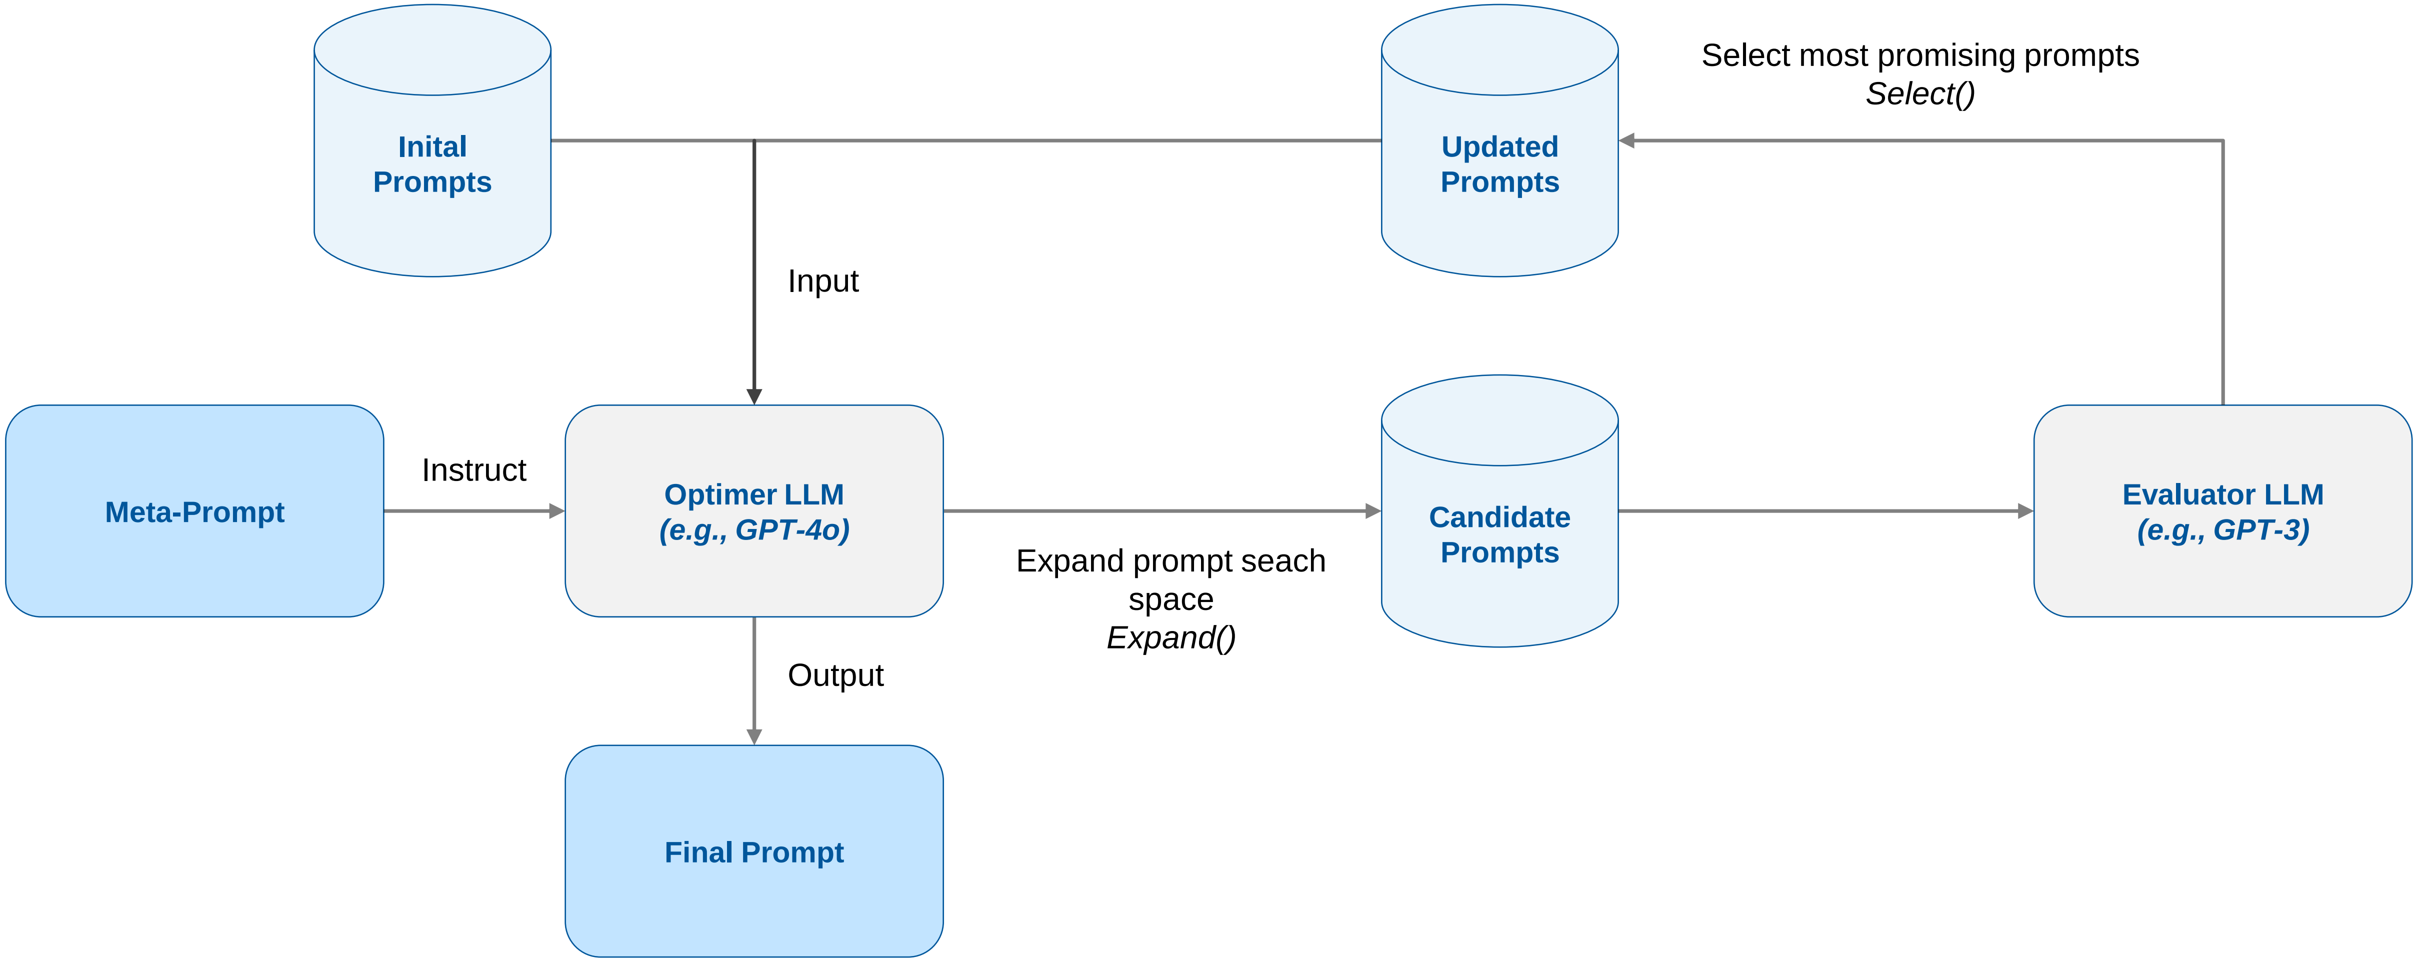

© sebis

3 classes of operators for incremental prompt refinement

RQ 1 – Initial Results

	Reflection-based	Evolution-based	Black-Box Optimization
Expansion	see figure below	see prompts below	see figure below
Selection	• Beam Search • MCTS • …	• Retain top k best performing prompts in population (fitness) • …	• Retain only one prompt in each iteration

Reflection-based

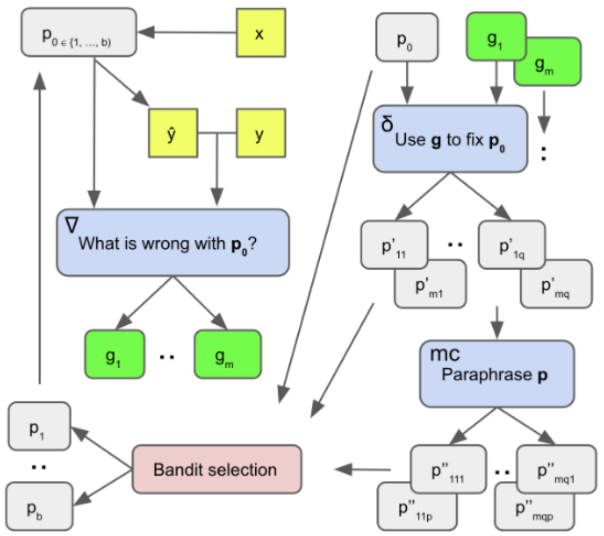

Evolution-based

Prompt 1: You are a categorizer, your mission is to ascertain the sentiment of the provided text, either favorable or unfavourable

Prompt 2: Assign a sentiment label to the given sentence from [negative, positive]

Crossover

Your mission is to ascertain the sentiment of the provided text and assign a sentiment label from [negative, positive]

Mutation

Determine the sentiment of the given sentence and assign a label from [negative, positive]

Black-Box Optimization

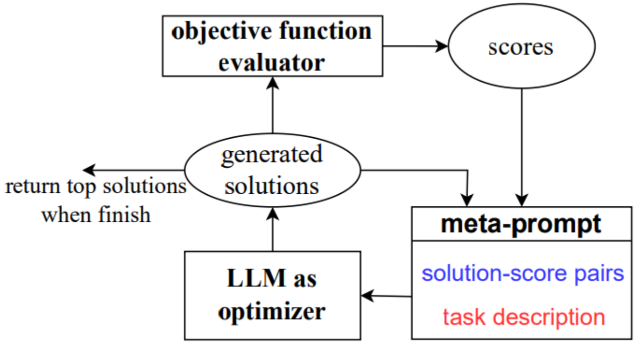

With Research Question 2 we want to find out how we can not only optimize the instruction but also the context, i.e., the “Prüfanleitungen” of the prompt

RQ 2 – Methodology

Research Question 2

How can these methods be adapted to optimize prompts, augmented with retrieved legal expert knowledge, for classifying the legality of German employment contracts without altering the semantics of the retrieved content?

Research Question 2.1

What are the syntactic dimensions to differentiate semantically equivalent versions of the “Prüfungsanleitungen” and how can they be represented?

Focus on next slide

Boundaries of the context representation search space

Research Question 2.2

How can the LLM-based APO methods be adapted s.t. not only the global instruction but also the local context as “Prüfungsanleitungen” is tuned while retaining semantic consistency?

APO Method development to
optimize entire prompt with pre-tuned Prüfanleitungen

© sebis

RQ 2.1 – 3 + 1 representations exist for formulating a Prüfungsanleitungen without altering semantics

RQ 2.1 – Methodology

Exception
Es besteht **<u>keine Verpflichtung</u>** zu [x], es sei denn mindestens eine der folgenden Ausnahmen ist erfüllt:
[Ausnahme 1]
…
[Ausnahme k]

Condition
Es besteht **<u>eine Verpflichtung</u>** zu [x], wenn mindestens eine der folgenden Bedingungen erfüllt ist:
[Bedingung 1]
…
[Bedingung k]

Reference
Die Pflicht zu [x] ist in [y] geregelt

© sebis

With Research Question 2 we want to find out how we can not only optimize the instruction but also the context, i.e., the “Prüfanleitungen” of the prompt

RQ 2 – Methodology

Research Question 2

How can these methods be adapted to optimize prompts, augmented with retrieved legal expert knowledge, for classifying the legality of German employment contracts without altering the semantics of the retrieved content?

Research Question 2.1

What are the syntactic dimensions to differentiate semantically equivalent versions of the “Prüfungsanleitungen” and how do they affect the performance of the LLM on the given downstream task?

Preliminary Experiment to
find optimal way (configuration) of formulating Prüfanleitungen

Research Question 2.2

How can the LLM-based APO methods be adapted s.t. not only the global instruction but also the local context as “Prüfungsanleitungen” is tuned while retaining semantic consistency?

Focus on next slide

Global APO Method development to
optimize entire prompt with pre-tuned Prüfanleitungen

RQ 2.2: Methodology – Algorithms for Context Optimization via Bandit Selection

RQ 2 – Methodology

Assuming no instruction optimization for simplicity

Components of APO method

Instruction optimization
Find the instruction that yields the highest accuracy

Focus

Context selection
Find the context (Prüfungsanleitungen) for a given clause that maximimizes the expectation for a correct classification

Focus

Context optimization
Find the best representation for the respective clauses that maximimizes the expectation for a correct classification

Pseudocode

```
For all clauses c in Clauses:
    prüfanleitungen = select_prüfanleitungen(c)

    reformulated_prüfanleitungen = []

    For all prüfanleitungen p in prüfanleitungen:
        r = select_representation(p)
        pr = reformulate_prüfanleitung(p, r)
        reformulated_prüfanleitungen.append(pr)

    y_hat = LLM(I, c, pr_1, ..., pr_n)

    reward = 1 if y == y_hat else 0

    For all prüfanleitungen p in prüfanleitungen:
        update_belief(p, r, reward)
```

RQ 2.2: Methodology – Algorithm for Context Optimization via Bandit Selection

RQ 2 – Methodology

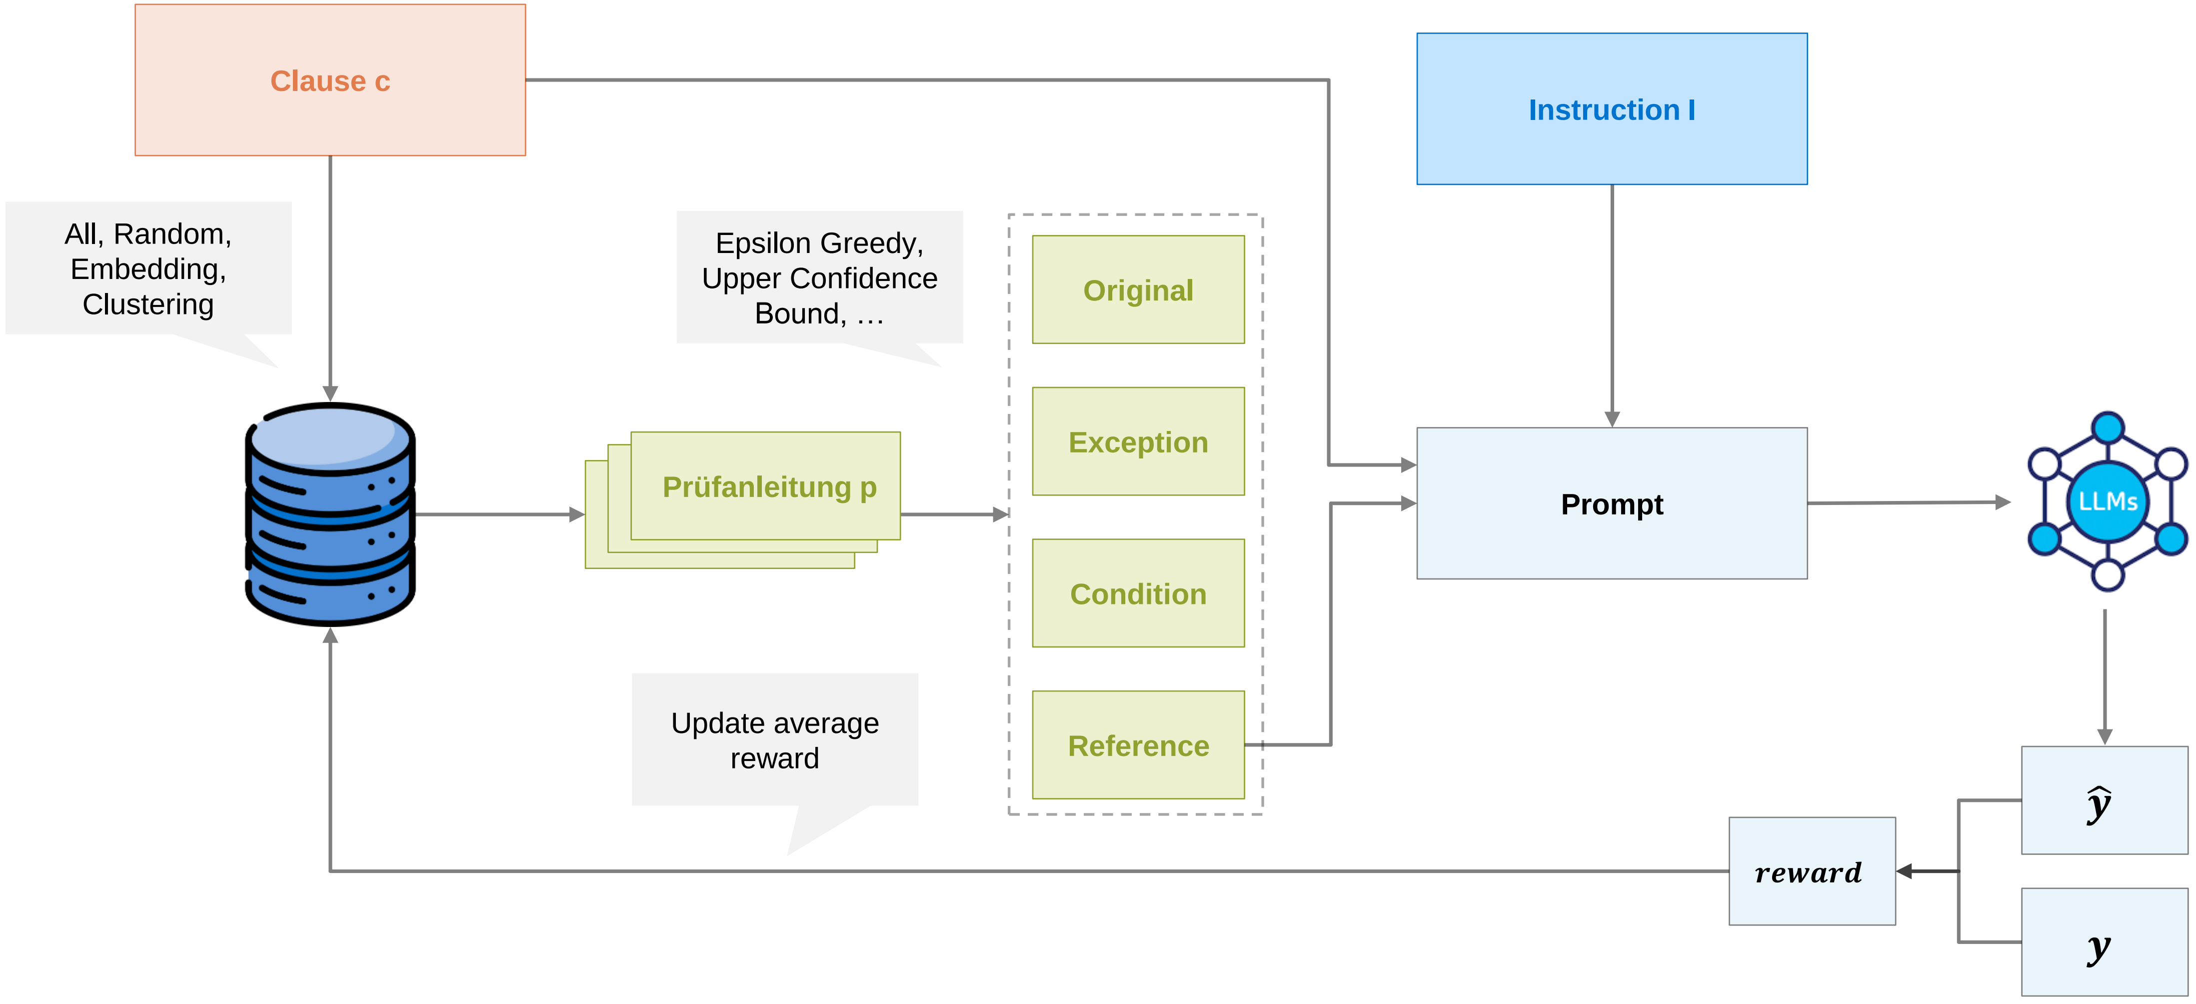

© sebis

RQ 3: Methodology

RQ 3 – Methodology

Perform experiments by measuring task performance for clause classification with Sylevenstein dataset

Context Selection

- No Prüfungsanleitungen
- All Prüfungsanleitung
- Random
- Embedding
- Clustering

Context Optimization

- All original
- Random
- Epsilon-Greedy
- UCB

Instruction Optimization

- No IO
- IO
- IO with clustering

© sebis

Outline

© sebis

Timeline

Preliminary timeline assumes registration on 26.08.24

Preliminary Timeline

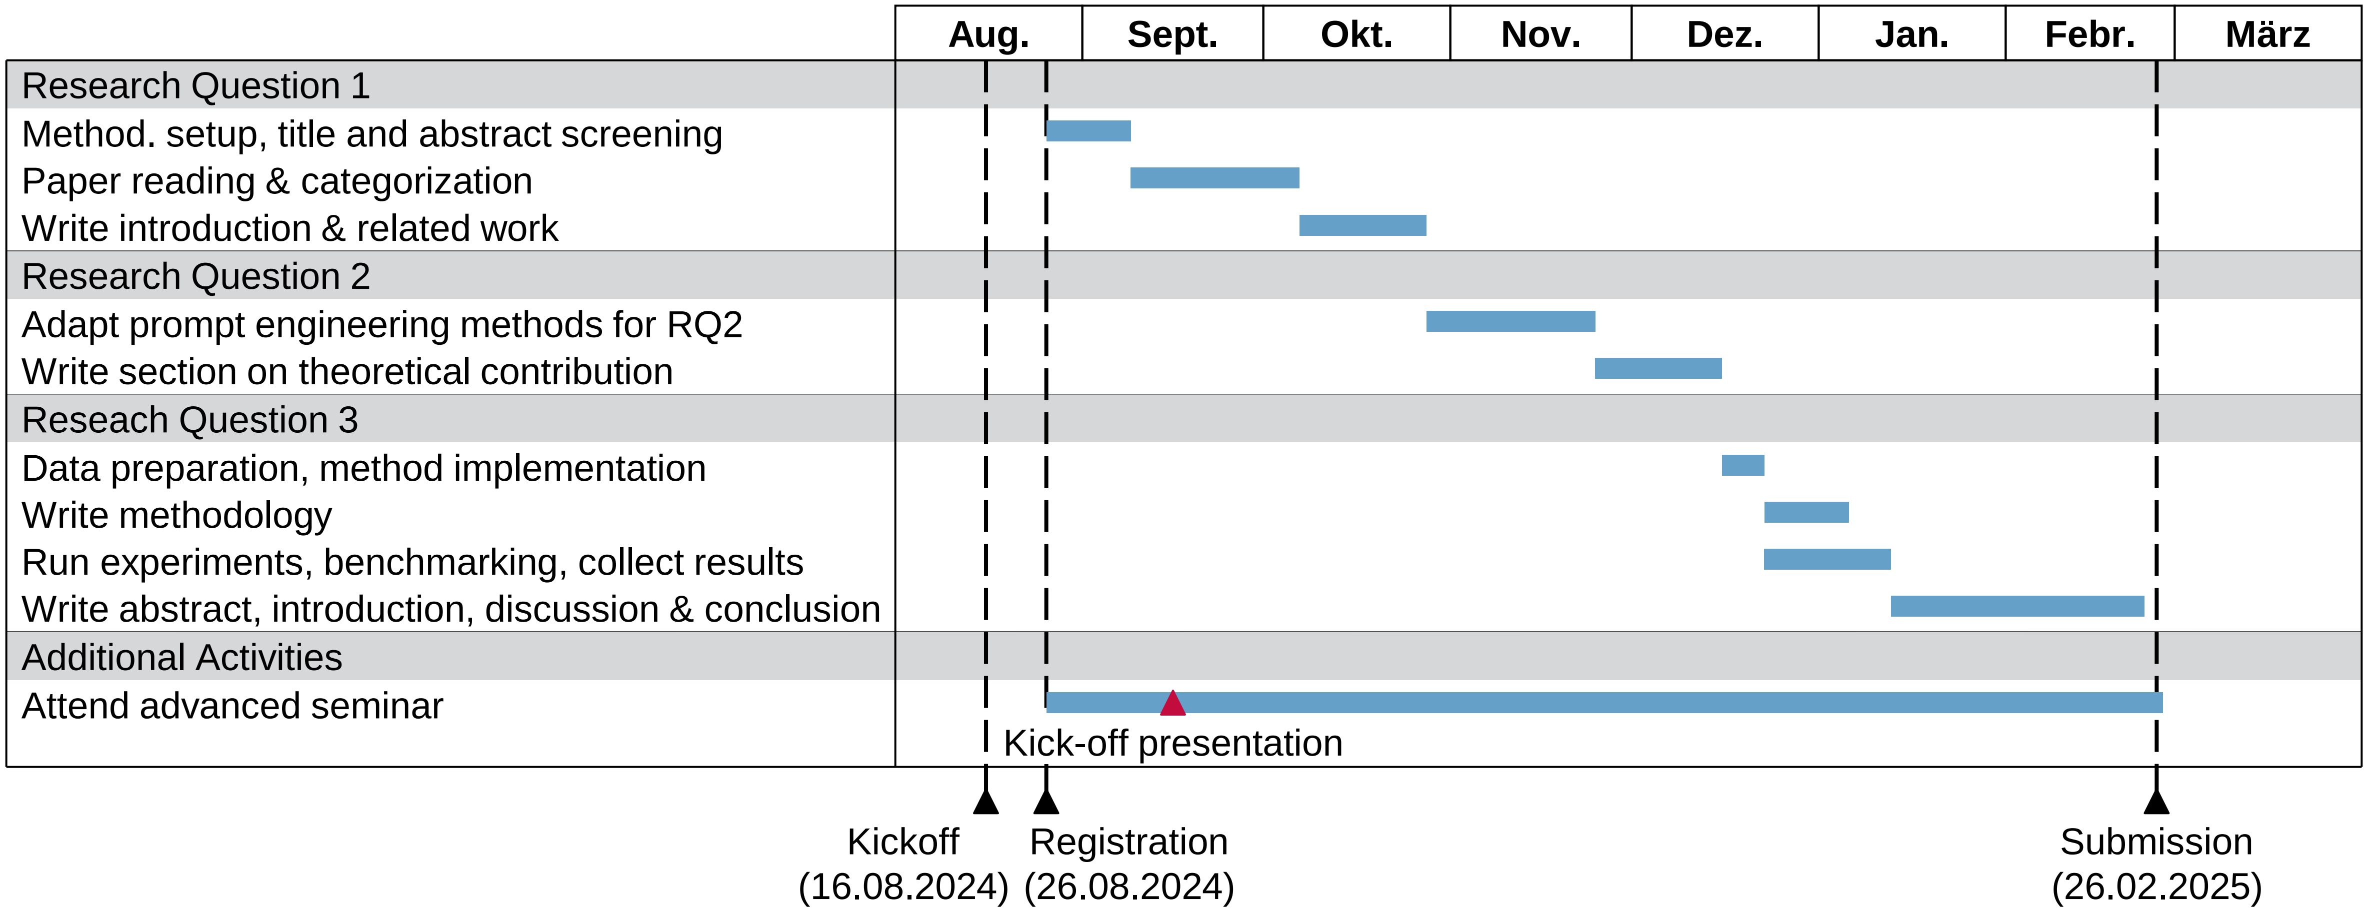

© sebis

Prof. Dr.

Florian Matthes

Technical University of Munich (TUM)
TUM School of CIT
Department of Computer Science (CS)
Chair of Software Engineering for Business Information Systems (sebis)

Boltzmannstraße 3
85748 Garching bei München

+49.89.289.17132
matthes@in.tum.de
wwwmatthes.in.tum.de

Sources

[1] A. Radford *et al.*, ‘Language models are unsupervised multitask learners’, *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.

[2] T. B. Brown, ‘Language models are few-shot learners’, *arXiv preprint ArXiv:2005. 14165*, 2020.

[3] J. Wei *et al.*, ‘Chain-of-thought prompting elicits reasoning in large language models’, *Advances in neural information processing systems*, vol. 35, pp. 24824–24837, 2022.

[4] X. Wang *et al.*, ‘Self-consistency improves chain of thought reasoning in language models’, *arXiv preprint arXiv:2203. 11171*, 2022.

[5] B. Lester, R. Al-Rfou, and N. Constant, ‘The power of scale for parameter-efficient prompt tuning’, *arXiv preprint arXiv:2104. 08691*, 2021.

[6] G. Qin and J. Eisner, ‘Learning how to ask: Querying LMs with mixtures of soft prompts’, *arXiv preprint arXiv:2104. 06599*, 2021.

[7] Y. Hao, Z. Chi, L. Dong, and F. Wei, ‘Optimizing prompts for text-to-image generation’, *Advances in Neural Information Processing Systems*, vol. 36, 2024.

[8] Y. Zhou *et al.*, ‘Large language models are human-level prompt engineers’, *arXiv preprint arXiv:2211. 01910*, 2022.

[9] T. Shin, Y. Razeghi, R. L. Logan IV, E. Wallace, and S. Singh, ‘Autoprompt: Eliciting knowledge from language models with automatically generated prompts’, *arXiv preprint arXiv:2010. 15980*, 2020.

[10] T. Zhang, X. Wang, D. Zhou, D. Schuurmans, and J. E. Gonzalez, ‘Tempera: Test-time prompting via reinforcement learning’, *arXiv preprint arXiv:2211. 11890*, 2022.

[11] M. Deng *et al.*, ‘Rlprompt: Optimizing discrete text prompts with reinforcement learning’, *arXiv preprint arXiv:2205. 12548*, 2022.

[12] A. Prasad, P. Hase, X. Zhou, and M. Bansal, ‘Grips: Gradient-free, edit-based instruction search for prompting large language models’, *arXiv preprint arXiv:2203. 07281*, 2022.

[13] Y. Guo, Y. Liang, C. Wu, W. Wu, D. Zhao, and N. Duan, ‘Learning to Plan with Natural Language’, *arXiv preprint arXiv:2304. 10464*, 2023.

[14] R. Pryzant, D. Iter, J. Li, Y. T. Lee, C. Zhu, and M. Zeng, ‘Automatic prompt optimization with" gradient descent" and beam search’, *arXiv preprint arXiv:2305. 03495*, 2023.

Prof. Dr.
Florian Matthes

Technical University of Munich (TUM)
TUM School of CIT
Department of Computer Science (CS)
Chair of Software Engineering for Business
Information Systems (sebis)

Boltzmannstraße 3
85748 Garching bei München

+49.89.289.17132
matthes@in.tum.de
wwwmatthes.in.tum.de

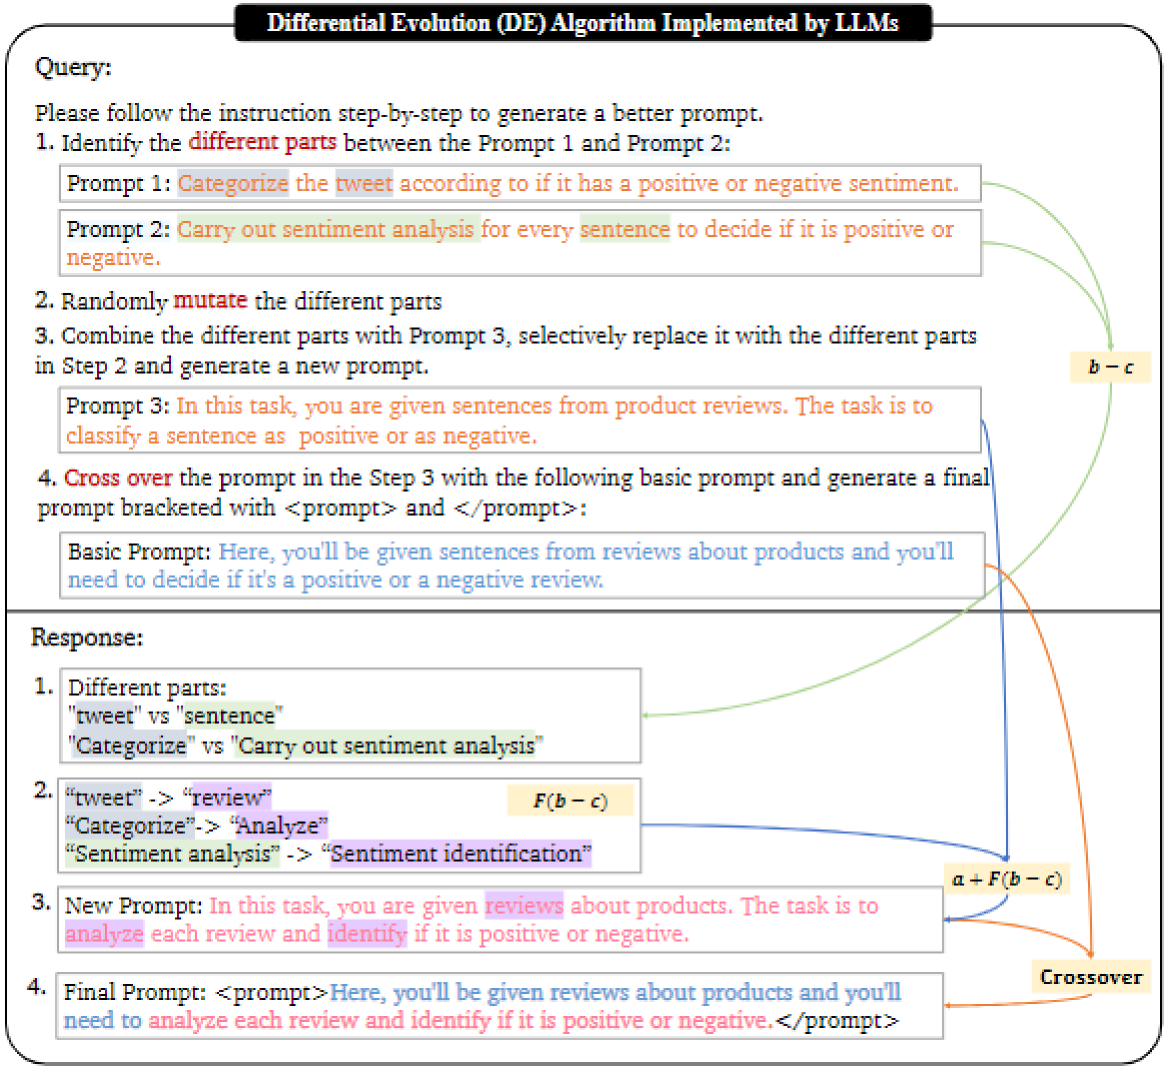

Figure 2: DE process implemented by LLMs (Evo($\cdot$) in Algorithm 1). In Step 1, LLMs find the different parts (words in ■ and ■) between Prompt 1 and Prompt 2 ($\mathbf{b} - \mathbf{c}$ in typical DE). In Step 2, LLMs perform *mutation* (words in ■) on them (imitation of $\mathbf{F}(\mathbf{b} - \mathbf{c})$). Next, LLMs incorporate the current best prompt as Prompt 3 with the mutated results in Step 2, to generate a new prompt (counterpart of $\mathbf{a} + \mathbf{F}(\mathbf{b} - \mathbf{c})$ in DE). Finally, LLMs perform *crossover* upon the current basic prompt p_i and the generated prompt in Step 3. See Figure 5 in Appendix B.2 for the complete response.