

# Automatic Prompt Optimization for Knowledge Graph Construction: Insights from an Empirical Study

Nandana Mihindukulasooriya  
IBM Research  
New York, USA  
nandana@ibm.com

Faisal Chowdhury  
IBM Research  
New York, USA  
mchowdh@us.ibm.com

Niharika S. D’Souza  
IBM Research  
San Jose, USA  
Niharika.DSouza@ibm.com

Horst Samulowitz  
IBM Research  
New York, USA  
samulowitz@us.ibm.com

## ABSTRACT

A knowledge graph (KG) represents a network of entities and illustrates relationships between them. KGs are used for various applications, including semantic search and discovery, reasoning, decision-making, natural language processing, machine learning, and recommendation systems. Automatic KG construction from text is an active research area. Triple (subject-relation-object) extraction from text is the fundamental building block of KG construction and has been widely studied, for example, in early benchmarks such as ACE 2002<sup>1</sup> to more recent ones, such as WebNLG 2020<sup>2</sup>, REBEL and SynthIE. In recent years, a number of works have explored the use of Large Language Models (LLMs) for KG construction. However, handcrafting reasonable task-specific prompts for LLMs is a labour-intensive exercise and can be brittle due to subtle changes in the LLM models employed. Recent work in NLP tasks (e.g. autonomy generation) uses automatic prompt optimization/engineering to address this challenge by generating optimal or near-optimal task-specific prompts given input-output examples.

This empirical study explores the application of automatic prompt optimization for the triple extraction task using experimental benchmarking. We evaluate different settings by changing (a) the prompting strategy, (b) the LLM being used for prompt optimization and task execution, (c) the number of canonical relations in the schema (schema complexity), (d) the length and diversity of input text, (e) the metric used to drive the prompt optimization, and (f) the dataset being used for training and testing. We evaluate three different automatic prompt optimizers, namely, DSPy, APE, and TextGrad and use two different triple extraction datasets, SynthIE and REBEL. Through rigorous empirical evaluation, our main contribution highlights that automatic prompt optimization techniques can generate reasonable prompts similar to humans for triple extraction. In turn, these optimized prompts achieve improved results, particularly with increasing schema complexity and text size.

## VLDB Workshop Reference Format:

Nandana Mihindukulasooriya, Niharika S. D’Souza, Faisal Chowdhury, and Horst Samulowitz. Automatic Prompt Optimization for Knowledge Graph Construction: Insights from an Empirical Study. VLDB 2025 Workshop: LLM+Graph.

## 1 INTRODUCTION

Knowledge graph (KG) construction is a widely studied research area due to the importance and usage of KGs in a wide range of applications. The phrase “knowledge graph” has been used in the literature since at least 1972 [13] and, to date, there exist several hundred methods for KG construction [46].

KGs can be constructed either from structured or semi-structured data using mappings such as RDB2RDF [35] or RML [7] or using unstructured data using various information extraction techniques such as Named Entity Recognition (NER), Relation Extraction (RE), and Open Information Extraction (OIE). In the public domain, crowdsourcing has also been used to build large KGs, such as Wikidata [40]. Nevertheless, crowdsourcing is not a feasible option for most industrial KGs where the knowledge is organization-specific and relevant text corpora are often confidential.

Recent advances in large language models, which have significantly improved performance on core NLP tasks, have transformed the research area of KG construction from text [17, 39]. Most of the early large language model (LLM)-based methods approach the problem in two key ways. The first is as an extraction of factual and common-sense knowledge using pre-trained language models via prompt-based prediction of masked objects from partial sentences describing complete triples [29]. The second is by fine-tuning an LLM on the autoregressive generation of text for triple extraction [23, 33]. Similarly, in-context learning with decoder-only transformer models is applied for relation extraction [18]. Alternatives such as Re2G [10] and GraphRAG [8] use LLMs to build Retrieval Augmented Generation (RAG) systems [19].

### 1.1 Prompt Engineering: Challenges and Opportunities

LLMs typically have a large number of parameters and are (pre)-trained on larger corpora, followed by iterative refinement using

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing [info@vldb.org](mailto:info@vldb.org). Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment. ISSN 2150-8097.

<sup>1</sup><https://www.ldc.upenn.edu/collaborations/past-projects/ace>

<sup>2</sup><https://synalp.gitlabpages.inria.fr/webnlg-challenge/challenge2020/>

Reinforcement Learning from Human Feedback (RLHF) [36] or Reasoning-oriented Reinforcement Learning [11]. These models are growing in popularity due to their capabilities in comprehending detailed instructions for diverse and complex downstream tasks [41] in a broad sense.

More often than not, the downstream performance of LLMs is heavily coupled with the quality of the prompt used to instruct the model. Studies have shown that large language models (LLMs) are quite sensitive to minor variations in prompt phrasing. For example, the addition, removal, or reordering of just a few tokens can lead to significant differences in task performance [21, 48]. Generic prompts do not typically produce good responses and the most effective prompts are almost always handcrafted by humans. This makes prompt curation a labour-intensive iterative process involving a substantial amount of manual experimentation. This process of optimizing the prompt language to elicit the best possible performance is referred to as “prompt engineering”. Given the human effort involved in the practice, prompt-engineering techniques are often brittle, non-transferable, and suffer from scalability issues [1]. Currently, prompt engineering is more of an art than a science, as a delicate balance is required in the design process to ensure clarity and specificity in the prompts, avoid ambiguity, and steer appropriate behaviour to obtain the desired output. More importantly, the human labor involved needs to be repeated whenever the underlying LLM is changed, since previously optimized prompts may no longer yield optimal results. For example, this applies when we switch to a different model, to a different parameter-size variant of the same model, or upgrade to new versions of the LLM trained using different strategies/additional data. This is further complicated when the prompt has to take into account a domain-specificity or a different language (direct translations of prompts often produce poor results [26])

## 1.2 Automatic Prompt Optimization: Motivation

Recent advances and superior pre-training strategies have resulted in LLMs that have shown remarkable, often superhuman, capabilities to generate reasonable responses even with limited or no examples. This has inspired a few works to explore the capabilities of an LLM in designing the best possible prompts for specific tasks. This process is referred to as automatic prompt optimization [32] or automatic prompt engineering [20]. The core idea is that the LLM is tasked with generating task-specific instructions where the task is presented via output demonstrations (with some given examples). It generates several instruction candidates, either via direct inference or a recursive process driven by scoring metrics. The LLM executes these instructions, and the best instruction improving the scoring metric is retained. Automatic prompt optimization can be applied to any task that is solved by prompting LLMs.

In this paper, we explore three prompt optimization approaches, namely, DSPy, TextGrad, and APE (detailed descriptions in Section 3), adapted to triple extraction. We extensively evaluate different settings guiding this task. Specifically, we explore the effect of changing (a) the prompting strategy, (b) the LLM being used for prompt optimization and task-execution, (c) the number of canonical relations allowed in the schema (i.e. the schema complexity),

(d) the length and diversity of the input text, (e) the metric driving the optimization, and (f) the dataset being used for training and testing.

We demonstrate that automatic prompt optimization techniques for triple extraction can generate reasonable prompts akin to human-generated prompts and thus achieve improved results. Specifically, the most significant gains were observed as the text size, context length, and schema complexity increases. To the best of our knowledge, this is the first work that analyses the impact of automatic prompt optimization for triple extraction for KG construction.

## 2 TRIPLE EXTRACTION FOR KNOWLEDGE GRAPH CONSTRUCTION

To build a Knowledge Graph from text corpora, a system needs to convert the unstructured natural-language text into a set of elemental triples. Triple extraction involves sub-tasks such as identifying the core entities mentioned in the text (Named Entity Recognition or NER), along with identification/extraction of relations that constitute the facts mentioned in the text (Relation Extraction or RE). These sub-tasks can be pipelined, by performing entity extraction first, and then extracting relations and triples or in a joint manner, where entities, relations, and triples are extracted simultaneously (end-to-end relation extraction). In this paper, we focus on end-to-end triple extraction, where the input is a natural language text and the expected output is a set of triples.

Triple extraction may be either in the form of closed relation extraction (where the set of relations is restricted to a set of known canonical relations) or of the open information extraction form [9]. Pre-curated canonical relations are useful when the resulting KG needs to be grounded by a pre-defined ontology. On the other hand, open information extraction requires relations to be spontaneously discovered from text and then represented using short descriptive phrases. This requires converting the phrases extracted from text to surface forms that are consistent/resolvable throughout the extraction process. Thus, when a large number of sparse and potentially diverse relations are involved, extensive post-processing may be needed. This includes clustering [45], canonicalization [6], and relation linking [34] to ground extracted relations to the canonical relations in the KG, which adds a layer of complexity to the KG-extraction process over the LLM-based processing. Since we are interested in evaluating the potential of LLMs for KG, we focus on closed triple extraction with a set of pre-known canonical relations.

In traditional triple extraction approaches, the set of canonical relations is used during the training phase to select the labels for relation classification or ranking. Invoking LLMs passes knowledge about these canonical relations to the model using either fine-tuning or prompt tuning. Instruction fine-tuning via InstructGPT [28], Reinforcement Learning from Human Feedback (RLHF) [3, 36] is known to assist LLMs in following a broad range of written instructions exactly while generating responses as Chain-of-Thought reasoning steps. On the other hand, in-context learning [25, 43] in conjunction with prompt engineering can teach models to perform new tasks by providing a few demonstrations of input-output pairs at inference.

To employ in-context learning or Chain-Of-Thought with a canonical set of relations during prompting, we follow an approach

Input:

**Text:** *Gabriel Maura was a member of the Real Academia de la Historia and the Royal Spanish Academy. He was also the relative of Germán Gamazo and the father of Antonio Maura. He was based in Madrid.*

**Allowed Relations:** ['academic degree', 'relative', 'launch contractor', 'organizer', 'operating area', 'proxy', 'venous drainage', 'official color or colors', 'has natural reservoir', 'mother house', 'member of', 'from narrative universe', 'next crossing downstream', 'Roman praenomen', 'student of', 'presynaptic connection', 'type foundry', 'father', 'invasive to', ...  
N other relations are omitted for brevity]

Output:

**Entities:** ['Gabriel Maura', 'Real Academia de la Historia', 'Royal Spanish Academy', 'Germán Gamazo', 'Antonio Maura', 'Madrid']

**Relations:** ['member of', 'relative', 'father', 'work location']

**Triples:**

(Gabriel Maura, member of, Real Academia de la Historia),  
(Gabriel Maura, member of, Royal Spanish Academy),  
(Gabriel Maura, relative, Germán Gamazo),  
(Gabriel Maura, father, Antonio Maura),  
(Gabriel Maura, work location, Madrid)

**Figure 1: An example of the triple extraction task with inputs and expected outputs (from the SynthIE dataset).**

analogous to Text2KGBench [24] and provide the schema information as input in the prompt. Overall, we formulate our task (illustrated in Figure 1) as follows: Given an input text written in natural language form and a list of allowed canonical relations, the LLM is asked to extract triples from the text in the form of (subject, relation, object). In addition to the triples, the LLM may be explicitly asked to generate entities and relations as an intermediate output.

### 3 AUTOMATED PROMPT OPTIMIZATION

Recent progress in Natural Language Processing has been exponential due to the availability of powerful LLMs such as GPT-4, Llama, and Claude. Users interact with models by supplying carefully handcrafted prompts, with the task formulated as a token completion problem. As alluded to before, prompt design can be a laborious exercise heavily reliant on domain expertise and ad hoc heuristic evaluation. Automated Prompt Optimization (APO) is a body of work that seeks to address this problem by leveraging systematic data-driven approaches.

The works of [5] provide a comprehensive taxonomy of automatic prompt optimization frameworks which refine prompts with no or minimal human intervention. Broadly, these methods can be categorized along multiple dimensions. These include: optimization space (i.e. discrete text-based vs. soft prompting or gradient-based), optimization targets (i.e. instructions vs. examples), optimization objective (i.e. task performance, safety, or generalizability), operators used to generate new prompts (e.g. purely model-based vs. iterative

refinement of example prompts), and iterative search strategies (e.g. Evolutionary vs. Monte Carlo search). Below, we provide a non-exhaustive overview of recent advances in APO to give readers a taste of this rapidly evolving field.

Ma et al. [22] uses a strategy based on the greedy search. Gradient-based approaches such as ProTeGI [30], MAPO [4], TextGrad [44], etc. use gradient descent-like algorithms to optimize prompt embeddings according to a predefined performance objective [42]. Starting from a human prompt, these methods learn continuous valued embeddings or prompt proposals using the gradient directions. On the other hand, approaches such as InstructZero [2] optimize low-dimensional “soft prompts” using LLMs instead of directly optimizing discrete instructions. They balance exploitation vs exploration in this space using Bayesian Optimization techniques. In terms of discrete optimization, approaches such as EvoPrompt [12] circumvent access to gradients and internal parameters of LLMs using Evolutionary Algorithms to generate candidate prompts, thus making them more general. Another approach to the same problem leverages in-context learning. For example, Automated Prompt Engineer (APE) [47] generates prompts based on a few real-life examples, tailoring them to the task without needing extensive training data. APE treats the instruction as a program and optimizes it by searching through a pool of instruction candidates proposed by an LLM to maximize a chosen score function. A complementary approach to APO uses reinforcement learning-based strategies such as OIRL [37], which model the interaction between the query-prompt pair via a reward model for proposing and evaluating candidate prompts suited for arithmetic reasoning tasks. In contrast, meta-Prompting [38] uses structural and syntactical aspects of tasks to create general prompts that guide the generation of task-specific prompts. In lieu of template-based approaches, DSPy [16] frames prompt engineering as a declarative, compiler-driven optimization task. DSPy poses prompting as a computational graph pipeline for text transformation, where LLMs are invoked through parameterized self-improving modules. This provides great flexibility across diverse tasks given a computational budget.

APO techniques have been demonstrated across a wide variety of NLP tasks [47] such as text classification, question answering, zero-shot learning, text generation, summarization, sentiment analysis, paraphrasing, code generation, and interactive dialogue systems. Going one step further, knowledge graph construction is an essential but complex NLP task with great relevance to data lakehouses because it enhances the usability, accessibility, and interoperability of the data stored within them. While KG-construction could greatly benefit from APO techniques, its application for triple extraction from text has not been studied before to the best of our knowledge.

## 4 EXPERIMENT BENCHMARKING SETUP

### 4.1 Task

Let  $\mathbf{x} = (x_1, x_2, \dots, x_n)$  denote a passage of natural language text with a list of terms, and let  $\mathcal{R} = \{r_1, r_2, \dots, r_m\}$  be a predefined set of canonical relation types. The goal of the triple extraction task is to identify: (a) A set of entity mentions  $\mathcal{E} = \{e_1, e_2, \dots, e_k\}$  from the text, (b) a subset of relations  $r \in \mathcal{R}$  expressed in the text, and (c) a set of directed triples  $\{(e_i, r, e_j) \mid e_i, e_j \in \mathcal{E}, r \in \mathcal{R}\}$  that are

represents the semantic relationships between the entity pairs that are mentioned in text. An example of the task, along with inputs and outputs, is illustrated in Figure 1.

## 4.2 Datasets

We used SynthIE and REBEL datasets in our experiments.

**SynthIE** [15] is a relation extraction benchmark generated using subgraphs of Wikidata and a synthetic data generation strategy. We used the *synthie\_text* (generated using *text-davinci-003*) for training, validation and testing automatic prompt optimizers. The dataset contains 888 relation types.

**REBEL** [14] is a relation extraction benchmark generated using Wikipedia abstracts and aligned them to Wikidata triples using wikilinks to identify relation candidates that are further validated using a Natural Language Inference (NLI) model. The dataset contained 1079 relation types.

## 4.3 Automatic Prompt Optimizers

We used three approaches with distinct flavours of APO for triple extraction, detailed below:

**DSPy** [16] (Declarative Self-improving Python) is an open-source toolkit for modular prompt design as programming with structured and declarative natural-language modules. DSPy allows specifying input/output behaviour as a signature and selecting a module (e.g., predict, Chain-Of-Thought, ReAct) to assign a strategy for invoking the LLM. DSPy natives parse the LLM output based on the pre-defined signature. It allows for the optimization of prompts using a training set, which is analyzed to create a ‘sub-dataset’ and task descriptions, propose better prompts with LLM-generated instructions, and select effective few-shots examples. Specifically, we used MIPROv2 (Multiprompt Instruction Proposal Optimizer Version 2) [27] to optimize instructions and few-shot examples simultaneously as a joint search problem. In the rest of the paper, when DPSy is mentioned, it is referring to DSPy with MIPROv2. It first bootstraps a set of few-shot example candidates, proposes instructions grounded in different task dynamics, and finds an optimized combination of these joint candidates using Bayesian optimization. Figure 2 shows an example of an automatically optimized prompt using DSPy.

**APE** [47] APE does not require any initial prompt. Given some examples of input and expected output, it generates several task instruction candidates, executes them using the target model, and selects the most appropriate instruction based on evaluation scores. By default, APE uses GPT-3 (text-davinci-002 to be more specific). We replace it with Granite-3.2-8B-Instruct, given its open-source availability. We used the forward generation template settings and randomly picked 10 examples from the SynthIE validation data to provide as few shots, and selected the generated prompt with the highest score based on the validation set of SynthIE.

**TextGrad** [44] Unlike statistical optimization in machine learning, which numerically estimates gradients for backpropagation, TextGrad leverages natural language feedback to iteratively improve LLM-generated prompts, effectively operating on the computational graph connecting the instruction to the output. We adopt

*Simple baseline starting prompt:*

*(input text, allowed relations) → (entities, relations, triples)*

*Automatically generated prompt with DSPy:*

*Given an input text and a set of allowed relations, extract entities, identify relevant relations between these entities that match the allowed relations, and generate triples representing these relationships in a structured format. The task involves processing the input text to recognize entities and relations, filtering relations based on the predefined set of allowed relations, and outputting the identified entities, relevant relations, and corresponding triples. The goal is to accurately represent relationships within the input text, focusing on the allowed relations provided.*

*### selected few-shot examples from training data ###*

**Figure 2: An example of an automatically optimized prompt using DSPy with the Llama 3.3-70B model and the SynthIE dataset as training/validation data.**

the TextualGradientDescent optimizer from the official Pytorch implementation in [44] with pre-set hyperparameters for triple extraction. The initial prompt is a basic instruction describing the task, required output format, allowed canonical relations, and an optional one-shot example input and output. The evaluation and training engines are chosen to be the Llama3-70b model. Training is performed for a maximum of 5 epochs, batch size of 16, allowing for early stopping using the validation set. The loss function driving the optimization is 1-F1, thus optimizing for both recall and precision simultaneously. Of the three approaches tried, text grad is the most computationally expensive and slow to train, as it requires at least two API calls for each element of the training set per epoch (one for the forward pass and a second for textual gradient estimation)

## 4.4 Language Models

The automatic prompt optimizers use language models to generate the prompts. In this work, we have used the following open-source large language models for prompt optimization and performing the triple extraction task: (1) Deepseek V3 (2) Qwen2.5-72B (3) Meta Llama3.3-70B Instruct (4) Mistral-8x22B (5) Microsoft Phi-4 14B (6) IBM Granite3.2-8B Instruct

## 4.5 Evaluation Metrics

**Entity, Relation, Triple extraction (P, R, F1)** To evaluate entity, relation, and triple extraction, the extracted values are compared against the ground truth values to estimate the precision, recall, and F1 metrics. These metrics are calculated for each test-case, and the final reported values are the macro-averaged.

**Relation-Type Mean Accuracy** To analyze the results at the level of relation types, we compute the mean accuracy per relation type. It captures the average proportion of correctly identified relation instances for each specific relation type, computed against all instances of the given type.

## 5 EMPIRICAL ANALYSIS AND INSIGHTS

In this section, we analyze the performance improvements obtained by automatic prompt optimization for triple extraction from text, and how certain aspects related to the task, such as the length of the text or the complexity of the schema and the system, i.e. LLM being used drive the performance.

In scenarios under investigation in the following sections, default settings are maintained, as follows.

**Default settings:** Prompting strategy: Predict (E-R-T), LLM: Llama3.3-70B, Number of relations per test case: 100, Test dataset: SynthIE, Optimizer: DSPy, Optimization metric: Triple F1 score.

### 5.1 Prompting strategy

**RQ1:** How do different prompting strategies influence the performance improvements in triple extraction?

**Varied setting:** Prompting strategy.

We curated six prompting-strategy variants that uses distinct styles to obtain the expected output. This includes (a) Predict with In-Context Learning, (b) Chain-Of-Thought, and (c) pipelining subtask prompts for Extract, Critique, and Refine. The expected output may be one of (i) Triple (T) or (ii) Entity, Relations, and Triples (E-R-T). The six resulting variants are in Table 1. Each setting has a baseline prompt (see Table 1) and is run through the prompt optimization process which proposes a refined prompt.

Table 2 outlines the performance improvements in terms of macro Precision, Recall, and F1 obtained by prompt optimization. We observe that all prompting strategies result in a performance gain, as seen in Table 2. Of these, Chain-Of-Thought asked to solely extract triples obtains the best triple extraction F1 of 0.73. The highest overall performance improvement is seen in the case of the Extract, Critique, and Refine pipeline run on entities, relations, and triples i.e., ECR (E-R-T) output, with a +16% improvement in triple extraction F1.

Figure 3 shows the performance improvement at the relation-type level between the baseline prompt and the optimized prompt. It separates the number of relations that saw an increase in the mean accuracy (+ positive diff), a decrease in mean accuracy (- negative diff), and no change in mean accuracy before and after the prompt optimization process. The relation types that didn't have a change are further separated into the incorrect ones (i.e., mean accuracy 0) and the correct ones. It shows that in all prompting strategies, there are more relation types having a performance improvement (green) than a performance degradation (red) in Figure 3. It also suggests that the optimized prompts retain the majority of correct relations (grey) from the baseline prompts.

Similarly, in Figure 4, we analyze if all approaches find the same types of relations easier to extract (extract them correctly always) and harder to extract (extract them incorrectly always). Each heatmap shows the overlap of fully correct relation types (Upper Right, Green) and fully incorrect relation types (Lower Left, Red) among the different prompting strategies for the baseline prompts and optimized prompts. The fully correct Relation-Types refer to the Relation-Types with a mean accuracy of 1.0, and fully incorrect Relation-Types refer to relations with a mean accuracy of 0. The diagonal shows the total number of fully correct and incorrect Relation-Types for each approach. We observe high overlap

between different approaches; for example, between Predict (T) and CoT (T), there is a 95% (423/444) overlap in fully correct relations and a 88% (128/145) overlap in fully incorrect relation types.

**Table 1: Different prompting strategies for applying APO.**

APO	Description	Baseline Prompt
Predict (T)	Extracts triples from input text using a given list of allowed relations	(input text, allowed relations) → (triples)
CoT (T)	Asks the model to think step-by-step and provide a reasoning in addition to the triples.	(input text, allowed relations) → (reasoning, triples)
ECR (T)	A 3-step pipeline: (E)xtract, (C)ritique, and (R)efine triples with three distinct prompts. Each step passes its output to the next stage.	P1: (input text, allowed relations) → (triples) P2: (input text, allowed relations, extracted triples) → (triple critique) P3: (input text, allowed relations, extracted triples, triple critique) → (refined triples)
Predict (E-R-T)	Similar to Predict (T) but also explicitly extracts lists of entities and relations.	(input text, allowed relations) → (entities, relations, triples)
CoT (E-R-T)	Similar to CoT (T) but also explicitly extracts lists of entities and relations.	(input text, allowed relations) → (reasoning, entities, relations, triples)
ECR (E-R-T)	A 3-step pipeline (E)xtract, (C)ritique, and (R)efine similar to ECR (T) but also explicitly extracts lists of entities and relations in addition to the triples. All outputs of the previous step is passed as input to the next step.	P1: (input text, allowed relations) → (entities, relations, triples) P2: ... → (entity critique, relation critique, triple critique) P3: ... → (refined entities, refined relations, refined triples)

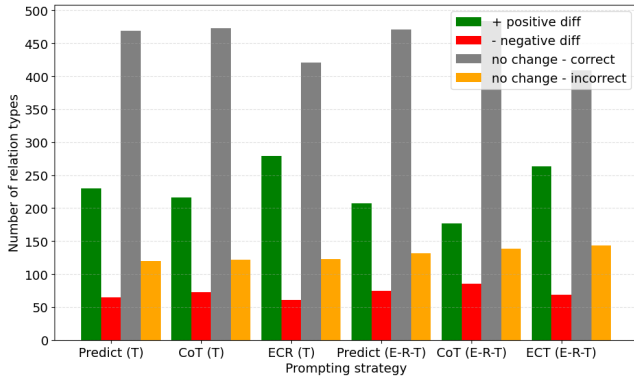
### 5.2 Prompt Generation LLM and Triple Extraction LLM

**RQ2:** How do different LLMs influence the performance improvements of automatic prompt optimisation techniques?

**Varied setting:** LLM used for prompt generation and testing.

**Table 2: Precision (P), Recall (R) , F1 metrics for Entity, Relation, Triple extractions using different prompting strategies. These experiments use Llama-3.3 70B as the LLM with 100 allowed relations per test case, evaluated on SynthIE small test set. White rows represent the baseline, and grey rows represent the optimized prompts.**

Prompting Strategy	Entity			Relation			Triple		
	P	R	F1	P	R	F1	P	R	F1
Predict (T)	-	-	-	-	-	-	0.69	0.60	0.63
	-	-	-	-	-	-	0.73	0.73	0.72
CoT (T)	-	-	-	-	-	-	0.71	0.60	0.64
	-	-	-	-	-	-	0.76	0.72	0.73
ECR (T)	-	-	-	-	-	-	0.60	0.59	0.59
	-	-	-	-	-	-	0.74	0.71	0.72
Predict (E-R-T)	0.98	0.90	0.93	0.80	0.66	0.71	0.68	0.58	0.62
	0.98	0.94	0.96	0.80	0.77	0.77	0.74	0.71	0.72
CoT (E-R-T)	0.98	0.91	0.94	0.81	0.68	0.72	0.69	0.59	0.63
	0.98	0.94	0.95	0.81	0.78	0.78	0.74	0.70	0.71
ECR (E-R-T)	0.95	0.94	0.94	0.62	0.65	0.62	0.52	0.56	0.53
	0.98	0.94	0.96	0.80	0.76	0.77	0.72	0.68	0.69



**Figure 3: Number of relations types with +ve / -ve / 0 differences in accuracy after prompt optimization.**

We performed a number of experiments using DSPy to probe the research question above and report results in Table 3. The length of the best prompts generated using Deepseek V3, Qwen2.5-72B, Llama3.3-70B, Mistral-8x22B, Phi-4 14B and Granite3.2-8B are 68, 81, 77, 58, 268 and 125 words long, respectively.

Note that the results achieved by all models are similar in performance for entity extraction, regardless of the exact model used to generate prompts. However, larger models (Deepseek V3, Qwen2.5-72B, and Llama3.3-70B) result in significantly better results for relation (and triple) extraction. Secondly, the best prompt generated using Llama3.3-70B is robust enough to obtain comparable results to those of the best prompt generated using other models, when testing with those. For example, with prompt generated using Llama3.3-70B and Phi-4 14B and tested using Phi-4 14B, we obtained

Predict (T)	✓ 468 ✗ 141	423	422	410	393	397
CoT (T)	128	✓ 444 ✗ 145	410	394	389	388
ECR (T)	125	122	✓ 454 ✗ 143	395	383	386
Predict (E-R-T)	119	120	116	✓ 442 ✗ 150	396	393
CoT (E-R-T)	121	124	119	132	✓ 420 ✗ 160	396
ECR (E-R-T)	123	126	122	133	151	✓ 423 ✗ 168
	Predict (T)	CoT (T)	ECR (T)	Predict (E-R-T)	CoT (E-R-T)	ECR (E-R-T)

**(4.a) Optimized Prompts**

Predict (T)	✓ 393 ✗ 189	354	294	357	347	305
CoT (T)	157	✓ 397 ✗ 188	296	353	350	305
ECR (T)	138	136	✓ 359 ✗ 220	301	297	271
Predict (E-R-T)	157	156	146	✓ 393 ✗ 190	364	316
CoT (E-R-T)	158	154	145	169	✓ 390 ✗ 196	326
ECR (E-R-T)	155	154	153	167	183	✓ 350 ✗ 225
	Predict (T)	CoT (T)	ECR (T)	Predict (E-R-T)	CoT (E-R-T)	ECR (E-R-T)

**(4.b) Baseline Prompts**

**Figure 4: Overlap of fully correct relation types, i.e., mean accuracy 1.0 (Upper, Green) and fully incorrect relation types, mean accuracy 0.0 (Lower, Red) of different baseline and optimized prompts. Diagonal shows the total count of fully correct (✓) and fully incorrect (✗) relation types for each optimized prompt.**

0.62 and 0.61 F1-scores respectively for triple extraction. Similarly, prompts generated with Granite3.2-8B and Llama3.3-70B obtain the same F1-scores of 0.54. It might imply that smaller models are capable of generating equally effective prompts compared to larger models. Furthermore, we observe that the choice of LLM during inference has a significant impact on the performance, as we obtain better results using larger models. This may imply that the choice of LLM at inference is more important than the choice of LLM that generates prompts.



**Table 3: Results for using different LLMs for automatic prompt generation and triple extraction tasks. Each model contains results with a baseline (no prompt optimization), a prompt generated using the same model, and a prompt generated from a different model. Models are ordered by parameter size.**

Validation /Test LLM	Prompt Gen. LLM	Entity			Relation			Triple		
		P	R	F1	P	R	F1	P	R	F1
Deepseek V3 671B	Baseline	0.96	0.96	0.95	0.82	0.72	0.75	0.68	0.60	0.63
	Deepseek V3	0.98	0.92	0.94	0.83	0.78	0.80	0.76	0.71	0.73
	Llama3.3-70B	0.98	0.95	0.96	0.84	0.81	0.82	0.78	0.75	0.76
Qwen2.5-72B	Baseline	0.97	0.95	0.96	0.80	0.72	0.75	0.71	0.64	0.66
	Qwen2.5-72B	0.98	0.94	0.96	0.79	0.77	0.77	0.73	0.70	0.71
	Llama3.3-70B	0.98	0.95	0.96	0.81	0.80	0.80	0.72	0.72	0.71
Llama3.3-70B	Baseline	0.98	0.90	0.93	0.77	0.68	0.71	0.68	0.59	0.62
	Llama3.3-70B	0.98	0.94	0.96	0.80	0.77	0.77	0.74	0.71	0.72
	Phi-4 14B	0.98	0.92	0.94	0.83	0.72	0.76	0.74	0.65	0.68
Mistral-8x22B	Baseline	0.96	0.96	0.95	0.64	0.61	0.61	0.55	0.52	0.53
	Mistral-8x22B	0.98	0.93	0.95	0.67	0.64	0.64	0.61	0.58	0.58
	Llama3.3-70B	0.98	0.94	0.95	0.69	0.68	0.67	0.59	0.58	0.58
Phi-4 14B	Baseline	0.95	0.97	0.95	0.66	0.55	0.59	0.54	0.48	0.50
	Phi-4 14B	0.98	0.92	0.95	0.73	0.67	0.69	0.65	0.59	0.61
	Llama3.3-70B	0.98	0.92	0.94	0.74	0.67	0.69	0.67	0.60	0.62
Granite3.2-8B	Baseline	0.99	0.83	0.89	0.55	0.44	0.47	0.47	0.37	0.40
	Granite3.2-8B	0.98	0.90	0.93	0.66	0.60	0.62	0.58	0.52	0.54
	Llama3.3-70B	0.98	0.90	0.93	0.67	0.59	0.61	0.59	0.52	0.54

### 5.3 Schema Complexity

**RQ3:** What is the impact of schema complexity (i.e., number of allowed relations) on relation extraction performance? Are optimised prompts more robust to changes in schema complexity?

**Varied setting:** Number of relations per test case.

To analyze this research question, we change the number of allowed relations that we include in the prompt. When a larger number of relations are included in the input text, we expect that the increased search space makes the task slightly harder for the systems with the increase in relation types. We examine eight settings with the number of allowed relation types ranging from 100 – 800.

Table 4 illustrates the impact on Precision, Recall, and F1 metrics for Entity, Relation, and Triple extraction. As seen, the performance of relation and triple extraction drops significantly with an increased number of allowed relations, suggesting an increase in task complexity. Figure 5 displays the performance difference between the baseline and optimized prompt at the level of test-case while varying allowed relations 100 to 800. Interestingly, the F1 score for entity extraction reduces with the optimized prompt. For instance, in the setting where there are 600 relations, it drops from 93% to 91%. Hence, entity extraction scores do not show the same trend as in relation and triple extraction.

**Table 4: Precision (P), Recall(R), F1 metrics for entity, relation, triple extractions across varying numbers of relation types per each test case. These experiments use LLAMA 3.3 70B as the LLM with Predict (E-R-T) as the prompting strategy, evaluated on the SynthIE small test set. White rows represent the baseline, and grey rows represent the optimised prompts.**


Relation Count Per Test Case	Entity			Relation			Triple		
	P	R	F1	P	R	F1	P	R	F1
rel 100	0.98	0.90	0.93	0.77	0.68	0.71	0.68	0.58	0.62
	0.98	0.94	0.96	0.80	0.77	0.77	0.74	0.71	0.72
rel 200	0.98	0.90	0.93	0.75	0.65	0.68	0.66	0.57	0.60
	0.98	0.94	0.96	0.77	0.73	0.74	0.71	0.67	0.68
rel 300	0.98	0.90	0.93	0.74	0.63	0.67	0.65	0.55	0.59
	0.98	0.94	0.95	0.77	0.72	0.74	0.71	0.66	0.68
rel 400	0.98	0.90	0.93	0.62	0.48	0.53	0.54	0.42	0.46
	0.97	0.93	0.95	0.69	0.64	0.66	0.63	0.59	0.60
rel 500	0.98	0.90	0.93	0.62	0.47	0.52	0.54	0.41	0.45
	0.97	0.93	0.94	0.69	0.64	0.65	0.63	0.58	0.60
rel 600	0.98	0.90	0.93	0.66	0.51	0.56	0.58	0.44	0.49
	0.92	0.92	0.91	0.69	0.66	0.66	0.62	0.58	0.60
rel 700	0.98	0.90	0.93	0.62	0.47	0.52	0.54	0.41	0.45
	0.96	0.92	0.94	0.67	0.62	0.63	0.61	0.56	0.58
rel 800	0.98	0.89	0.93	0.57	0.42	0.47	0.50	0.37	0.41
	0.95	0.92	0.93	0.67	0.62	0.63	0.61	0.57	0.58

### 5.4 Examining the effect of increased context length

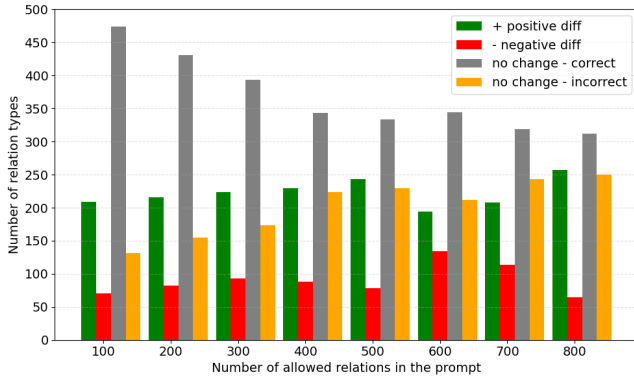
**RQ4:** How does APO perform at entity/relation/triple extraction as the size of the input text is increased?

**Varied setting:** Text length (1x, 5x, 10x test cases at a time).

In this experiment, we analyze the effect of increasing the context size of the input text on the performance of Automatic Prompt Optimization at entity/relation/triple extraction. Keeping the LLM



**Figure 5: Difference in Macro F1 Scores (Optimized - Baseline) at different number of allowed relations. F1 is taken as the macro average of all test cases.**



**Figure 6: Number of relation types with +ve / -ve / 0 differences in mean accuracy after prompt optimization at each relation count (100 to 800) setting.**

fixed as Llama 3.3-70b, we systematically increase the length of the text in each example (in our small test set of 1500 examples) by augmenting it with text from the larger test corpus (10000 - 1500 examples). In the first setting, text examples that are most similar to those in our small test set are selected and appended, referred to as Related Augmentation. The second setting, named Adversarial Augmentation, repeats the same, but augments the data with the most dissimilar candidates. In the third setting, passages are randomly selected for augmentation which we refer to as Random Augmentation. We examine increasing the size of the input text 5 and 10 fold and compare it with the results from Section 5.1. Text similarity is determined using a sentence transformer model. The extraction results are summarized in Table 5. We also indicate the distinct relation types in each experiment, a quantitative proxy for the difficulty of the extraction task. The adversarial Augmentation and Random Augmentation have noticeably higher numbers of distinct relation types as compared to the related text case.

As expected, the performance of the baseline and trained DSPy modules degrades as the context length increases. The degradation is most severe in the cases where the text is augmented adversarially, with relation and triple recall suffering the most. Entity extraction performance remains fairly consistent. Across the board, we observe that training for APO improves performance. We postulate that exposure to larger-sized text and especially their selection as few-shot examples, steers the optimization in a direction that helps better anticipate longer context. Improvements in performance are the largest in the case of Related Augmentation and the smallest in the Adversarial selection case. Finally, we notice that the performance in the Adversarial Augmentation case is noticeably worse than in the case of Random Augmentation, although the number of distinct relations is comparable. This suggests that the coherence of the text passage is an important factor influencing the difficulty of the extraction task.

## 5.5 Train, Validation and Test Datasets

**RQ5:** What is the impact of using train/validation splits from a different dataset for selecting few-shots and validating the generated prompts compared to using the train/validation splits from the

**Table 5: Quantifying the effect of increased context length on APO approaches for entity, relation and triple extraction. Three different settings are examined, where we merge (a) related text, (b) adversarially selected text, and (c) randomly selected text to each test example. White rows represent the baselines, and grey rows represent the optimized prompts. Numbers in the brackets indicate number of distinct relation types**

Text size (Rel. Types)	Entities			Relations			Triples		
	P	R	F1	P	R	F1	P	R	F1
<b>(a) Merging Related Text</b>									
1x (3.10)	0.98	0.90	0.93	0.77	0.68	0.71	0.68	0.59	0.62
	0.98	0.94	0.96	0.80	0.77	0.77	0.74	0.71	0.72
5x (10.25)	0.98	0.87	0.92	0.80	0.56	0.65	0.70	0.50	0.58
	0.99	0.92	0.95	0.77	0.67	0.71	0.71	0.63	0.66
10x (17.11)	0.99	0.86	0.91	0.82	0.47	0.59	0.70	0.43	0.53
	0.99	0.92	0.95	0.79	0.58	0.67	0.71	0.55	0.62
<b>(b) Merging Adversarially Selected Text (most unrelated)</b>									
1x (3.10)	0.98	0.90	0.93	0.77	0.68	0.71	0.68	0.59	0.62
	0.98	0.94	0.96	0.80	0.77	0.77	0.74	0.71	0.72
5x (13.57)	0.96	0.86	0.89	0.80	0.56	0.65	0.69	0.49	0.57
	0.99	0.85	0.91	0.83	0.58	0.67	0.72	0.51	0.59
10x (23.89)	0.96	0.85	0.89	0.83	0.48	0.60	0.72	0.43	0.54
	0.98	0.93	0.95	0.83	0.66	0.73	0.75	0.61	0.67
<b>(c) Merging Randomly Selected Text</b>									
1x (3.10)	0.98	0.90	0.93	0.77	0.68	0.71	0.68	0.59	0.62
	0.98	0.94	0.96	0.80	0.77	0.77	0.74	0.71	0.72
5x (14.35)	0.99	0.85	0.91	0.83	0.58	0.67	0.72	0.51	0.59
	0.98	0.92	0.95	0.82	0.68	0.74	0.74	0.62	0.67
10x (26.20)	0.98	0.84	0.90	0.85	0.51	0.63	0.73	0.45	0.55
	0.99	0.91	0.95	0.84	0.63	0.72	0.75	0.57	0.64

dataset as the test split?

**Varied setting:** Train/validation dataset, prompt strategy.

DSPy uses a training set to select the few-shot candidates and create a task/dataset summary that can be used as input to the LLM to derive prompt instruction candidates. In this experiment, we explore the use of different combinations of training/validation and testing datasets from distinct datasets and examine how the results change compared to using the same dataset for training and testing. We examine this behaviour using Predict (T) and Predict (E-R-T) prompting strategies.

In DSPy, few-shot candidates are selected from a training set and task/dataset summary is created. This is subsequently used as input for deriving prompt instruction candidates from an LLM. In this experiment, we explore the impact of using training/validation splits for different datasets compared to the test split. Specifically, we investigate how the results differ when the same dataset is used for both training and testing versus when different datasets are employed. We evaluate this behaviour using the Predict (T) and Predict (E-R-T) prompting strategies.

As shown in Table 6, the results demonstrate that using the same dataset for both few-shot selection and prompt validation



significantly outperforms the case when distinct datasets are used for training/validation and testing. For instance, when SynthIE is used as the test set, the Predict (E-R-T) prompt improves triple F1 by +0.01 when trained on the REBEL dataset, compared to a +0.08 improvement when the same dataset is used for both training and testing. Similarly, with Predict (T) and SynthIE as the test set, using REBEL for training/validation negatively impacts the optimized prompt, reducing triple F1 by -0.04 points. When REBEL is used as the test set, a similar trend is observed, with a larger margin of improvement when the same dataset is used for both training and testing compared to when different datasets are employed. Our hypothesis is that using few-shot examples and dataset summaries generated from a dataset that differs from the final test introduces some bias, leading to a prompt that is suboptimal for the test set.

**Table 6: Results for training and testing across two different triple extraction datasets.**

Test Set	Prompt Strategy	Train/Val Set	Entity F1	Rel. F1	Triple		
					P	R	F1
SynthIE	Predict (T)	N/A	-	-	0.69	0.60	0.63
		SynthIE	-	-	0.73	0.73	0.72
		REBEL	-	-	0.73	0.52	0.59
	Predict (E-R-T)	N/A	0.93	0.71	0.68	0.58	0.62
		SynthIE	0.95	0.77	0.72	0.70	0.70
		REBEL	0.91	0.71	0.71	0.58	0.63
REBEL	Predict (T)	N/A	-	-	0.24	0.29	0.23
		REBEL	-	-	0.37	0.35	0.33
		SynthIE	-	-	0.22	0.37	0.25
	Predict (E-R-T)	N/A	0.73	0.36	0.26	0.28	0.24
		REBEL	0.77	0.48	0.34	0.42	0.35
		SynthIE	0.74	0.39	0.24	0.35	0.26

## 5.6 Prompt Optimization Approach

**RQ6:** What prompt optimization approach works best for triple extraction?

**Varied setting:** Prompt optimization approach.

We experiment with three prompt optimization systems for triple extraction, namely, DSPy, TextGrad, and APE. As mentioned in Section 4.3, these methods differ in how they approach the problem. We include the best possible prompt (as generated by the corresponding system) for each case and then input it to the Llama3.3-70B model at test time. As seen in Table 7, the APE system is very competitive with the other methods despite being more simplistic in formulation and not requiring an initial human prompt as input. However, since APE needs to venture a guess about the task without any initial prompt and only a couple of input-output pairs, the optimal instruction is found to be generic (in a linguistic and descriptive sense) and not as task-specific compared with the two other approaches. For example, the APE system generated the following

**Baseline Prompt:** I have a friend an instruction. Based on the instruction, he produced the following input, output pairs. <Few Shot Examples>. Complete the following text. The instruction is

**Output:** Create a knowledge graph representation for the provided information, where each input is transformed into a list of dictionaries. Each dictionary should contain "subject", "object", and "predicate" as keys, with their respective values being the surface form of the subject/object and predicate phrase.

While the DSPy generated -

**Baseline Prompt:** (input text, allowed relations) → (entities,relations,triple)

**Output:** Given a text and a predefined set of allowed relations, extract the entities mentioned in the text and determine the relationships between them based on the allowed relations. Identify the subject, object, and relation for each extracted relationship and represent them as triples. Ensure that the extracted entities and relations are accurately matched to the allowed relations, and provide a comprehensive list of entities, relations, and triples that reflect the semantic relationships present in the input text.

And textgrad resulted in -

**Baseline Prompt:** Extract triples from the following text using ONLY the allowed relations. The output format should be (subject, object, relation). <EXAMPLES>. Do not include additional explanations or text.

**Output:** Extract concise and directly related knowledge graph triples from the text, focusing on explicit entity relationships, and only utilize the specified allowed relations to maximize precision and recall, ensuring that the subject and object are directly mentioned in the text and the relation is clearly implied.

According to the results, DSPy does have a small edge over the other two in performance.

**Table 7: Results for generating prompts with different automatic prompt optimizers.**

Automatic Prompt Optimizer	Entity			Relation			Triple		
	P	R	F1	P	R	F1	P	R	F1
Baseline	0.98	0.90	0.93	0.77	0.68	0.71	0.68	0.59	0.62
DSPy	0.98	0.94	0.96	0.80	0.77	0.77	0.74	0.71	0.72
APE	0.98	0.94	0.95	0.79	0.77	0.76	0.72	0.67	0.69
TextGrad	0.98	0.93	0.95	0.79	0.75	0.76	0.71	0.65	0.67

## 5.7 Optimization metric

**RQ7:** What is the impact of the optimization metric on the performance of the optimized prompt?

**Varied setting:** Prompt optimization metric.

Recall that DSPy selects the few-shot candidates from the training data and generates a set of instructions using an LLM. The candidate search space, including both the instructions and few-shot examples, is then evaluated using Bayesian optimization. The validation set is used to assess the candidates and guide the search strategy. During the process, a predefined optimization metric drives

the optimization process, i.e. to guide the candidate selection and evaluate against the validation set.

To explore the impact of the optimization metric, we vary the metric used during training to select the optimal prompts. We aim to determine whether this variation influences the prompt’s performance at test time as evaluated against the test set. Specifically, we seek to verify whether a prompt optimized for F1, Precision, or Recall performs better on the corresponding metric during testing.

Table 8 indicates the results for different settings with the three distinct optimization metrics. From the results, we observe that when triple precision is used as a metric, the triple precision on the test set improves slightly (0.75 compared to 0.71 and 0.72), as anticipated. Similarly, Recall also drives improvement in test Recall (0.71 compared to 0.66 and 0.70). Overall, this performance profile suggests that the optimization metric indeed influences the prompt’s behavior, but rather moderately in the specific setting evaluated. In addition, we manually inspect the three generated prompts to identify any linguistic signals correlating with the metrics. However, no notable instructional differences were found that drive the LLM in a particular direction.

**Table 8: Results for optimizing the prompt using different optimization metrics.**

Optimization Metric	Entity			Relation			Triple		
	P	R	F1	P	R	F1	P	R	F1
Baseline	0.98	0.90	0.93	0.80	0.66	0.71	0.68	0.58	0.62
Triple F1	0.98	0.94	0.95	0.79	0.77	0.77	0.72	0.70	0.70
Triple Precision	0.98	0.92	0.95	0.83	0.72	0.76	0.75	0.66	0.69
Triple Recall	0.98	0.95	0.96	0.78	0.79	0.77	0.71	0.71	0.70


## 5.8 Training Computation Cost

**RQ8:** What is the impact of increasing the number of LLM calls during optimization on the final performance?

**Varied setting:** APO Hyperparameters such as number of candidates, trials, mini-batch validation size, validation set size.

In this study, we define the computational cost of training as the number of LLM calls made during the process. In DSPy, the number of LLM calls is determined by several hyperparameters, including: the number of candidate prompts generated, the number of optimization trials conducted, the size of the mini-batch evaluated in each trial, the `full_eval_steps` (i.e., the number of mini-batch trials after which a full evaluation of the validation set is performed), and the size of the validation set. To assess the impact of these factors on the final performance of the generated prompt, we conducted experiments where we varied these hyperparameters.

As shown in Figure 7, both relation extraction and triple extraction F1 scores stabilise after a relatively small number of LLM calls. The stability of these results can be numerically quantified via the Coefficient of Variation (CV) in the triple test metrics (i.e. the relative ratio of standard deviation to mean), with lower values implying that the variation is less pronounced (lower sensitivity to hyperparameter choices). Increasing the search space, by increasing the number of candidates and trials ( $CV_{\text{TripleR}} = 0.02$ ,  $CV_{\text{TripleP}} = 0.026$ ,  $CV_{\text{TripleF1}} = 0.019$ ), or the validation set



**Figure 7: Training computation cost (in LLM calls) vs the relation and triple extraction F1.**

size ( $CV_{\text{TripleR}} = 0.014$ ,  $CV_{\text{TripleP}} = 0.014$ ,  $CV_{\text{TripleF1}} = 0.036$ ), does not result in a drastic change in performance. These findings suggest that beyond a certain point, expanding the search space and increasing computational effort do not yield proportional gains in performance.

## 6 QUALITATIVE ANALYSIS OF GENERATED PROMPTS

Three of the co-authors manually evaluated the top 22 prompts generated by Llama3.3-70B model using DSPy in different settings as described in Section 4. For each of these prompts, we asked ourselves “*Would a human prompt engineer write this prompt and consider it an optimal or near-optimal for the triple extraction task, i.e., do you expect a good result when this prompt is provided to an LLM?*”. We used for 4 categories to mark the prompts – Good, Short, Ok and Verbose. For these 22 prompt annotations, all three human annotators found that 3 prompts are Verbose and 8 prompts are Good. For the remaining 11, at least one (but not all 3) annotators marked 8 prompts as good.

The longest generated prompt marked “good” by all the annotators consists of 100 words -

Given a piece of text and a set of allowed relations, identify the entities present in the text, determine the relationships between these entities based on the allowed relations, and construct triples in the form of subject-relation-object representing the extracted information. The goal is to accurately extract entities, relations, and triples from the input text, utilising the provided allowed relations to guide the relationship identification process. Ensure the output includes a list of identified entities, a list of relations found between these entities, and a list of triples that effectively capture the relationships between the entities in the text.

Whereas the shortest prompt marked “good” by all the annotators consists of 77 words -

Given a text and a predefined set of allowed relations, extract the entities mentioned in the text and determine the relationships between them based on the allowed relations. Identify the subject, object, and relation for each extracted relationship and represent them as triples. Ensure that the extracted entities and relations are accurately

matched to the allowed relations, and provide a comprehensive list of entities, relations, and triples that reflect the semantic relationships present in the input text.

Interesting, both the prompts yielded similar results (when fed into the Llama3.3-70B) on the validation data - *F1-Entities*: 0.95/*F1-Relations*: 0.77/*F1-Triples*: 0.70.

## 7 CONCLUSIONS

In this empirical study, we explored the application of automatic prompt optimization for triple extraction via extensive experimental benchmarking. We evaluated different settings by changing (a) the prompting strategy, (b) the LLM being used for prompt optimization and task execution, (c) number of canonical relations in the schema (schema complexity), (d) the length and diversity of input text, (e) the metric used to drive the prompt optimization, and (f) the dataset being used for training and testing. Based on the empirical results, automatic prompt optimization is a viable strategy for generating reasonable prompts for the triple extraction task, providing robust improvements against baseline prompts. The three prompt optimization strategies examine, namely, DSPy, APE, and TextGrad, consistently outperformed the baseline prompts in precision, recall, and F1 metrics for entity, relation, and triple extraction. Additionally, qualitative analysis of the generated prompts showed that a majority of the generated prompts were reasonably interpretable and in alignment with human intuition. Specifically, they provide explicit details of the task, desired output format, and some guidance on how to perform the task.

The experiments with different prompting styles showed that both simple inference prompts (predict) as well as advanced prompting with a Chain-of-Thought (Cot) rationale or a pipelined Extract-Critique-Refine (ECR) prompts provided improvements during prompt optimizations. These observations suggest that the automatic optimizers can adapt and propose effective refinements across the board in scenarios related to triple extraction. Analysis at the relation-type level showed that optimized prompts tend to retain or improve performance on the majority of relation-types present in the dataset, without losing performance on those relation-types that were already handled correctly by the baseline. This indicates that APO not only boosts overall macro-level metrics but also does so broadly across different relation types, suggesting that the improvements are generalizable, i.e. do not narrowly focus on a few types.

Experiments in increased schema-complexity (i.e., the number of allowed relation types) and input context length (1x, 5x, and 10x of the original input text) showed that automatic prompt optimization was considerably more beneficial the more challenging the scenario. When the schema became more complex with a larger number of relations (from 100 up to 800 relations), extraction difficulty increased, and both baseline and optimized prompt performance drops in terms of relation and triple F1. Nevertheless, the decline was relatively more significant for baseline prompts, implying that the optimized prompts were more resilient to increased schema-complexity. Increasingly long texts (by concatenating multiple test phrases together), both baseline and optimized prompt performance deteriorated, especially on relation and triple extraction (with relation extraction recall dropping due to challenges in identifying all relevant relations in a large blob of text). Optimized prompts can

mitigate these context-length performance losses by still outperforming the baseline performance metrics for the single test case 1x all the way up to 10x. Descriptive optimized prompts with the right set of few-shot demonstrations seem to scale to larger input lengths quite gracefully.

**Limitations and Challenges.** When tested with different LLMs, it was observed that the triple extraction performance was mostly dominated by the capabilities of the task execution LLM rather than the LLM used to generate the prompts. The results imply that the choice of LLM at inference time seems more critical than the choice of LLM for prompt generation. In our experiments, prompts generated using a larger model (Llama 3.3-70B) were reasonably robust when applied to a different model (e.g. a 14B parameter model) on triple extraction. For example, the best prompt from Llama-70B tested on a smaller model achieved nearly the same triple F1 as the smaller model’s own optimized prompt. This suggests that there is some degree of “model-agnostic prompt quality” maintained in the optimised prompts (typically a good prompt containing precise instructions that are generally useful).

When tested with multiple datasets, our results also highlight some limitations in cross-dataset prompt transferability. For example, when we used a different dataset’s training examples to optimize prompts (this simulates a shift in domain/data distribution), the improvements reduced significantly compared to using the examples from the same dataset. In particular, an optimized prompt trained on a secondary dataset yielded only marginal gains (on  $\sim 1\%$  F1 increase) when tested on the target dataset’s text, compared to the much larger improvements ( $\sim 8\%$  or more) obtained when the training and testing data were from the same dataset. This highlights a set challenges for future work in prompt optimization.

**Future Work.** We plan on exploring (a) cross-lingual prompt [31] optimization, where we will apply automatic prompt optimization techniques to multilingual settings, and (b) interactive human-in-the-loop prompt optimization, i.e. combining automatic prompt optimization with human feedback loops. An interactive APO approach could allow a human analyst to guide the process by injecting domain knowledge or preferences (for example, selecting among automatically generated prompt candidates or refining them after a few iterations). Such approaches might fit well with agentic workflows that integrate tool-calling capabilities and decision making on when to get a human expert involved.

## REFERENCES

- [1] Simran Arora, Avani Narayan, Mayee F Chen, Laurel Orr, Neel Guha, Kush Bhatia, Ines Chami, and Christopher Re. 2023. Ask Me Anything: A simple strategy for prompting language models. In *The Eleventh International Conference on Learning Representations*.
- [2] Simran Arora, Avani Narayan, Mayee F. Chen, Laurel Orr, Neel Guha, Kush Bhatia, Ines Chami, Frederic Sala, and Christopher Ré. 2024. InstructZero: Efficient Instruction Optimization for Black-Box Large Language Models. *ICML 2024*.
- [3] Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems* 30 (2017).
- [4] Anthony Cui, Pranav Nandyalam, and Kevin Zhu. 2025. Introducing MAPO: Momentum-Aided Gradient Descent Prompt Optimization. *NAACL SRW 2025*.
- [5] Wendi Cui, Jiaxin Zhang, Zhuohang Li, Hao Sun, Damien Lopez, Kamalika Das, Bradley Malin, and Sricharan Kumar. 2025. Automatic Prompt Optimization via Heuristic Search: A Survey. *arXiv preprint arXiv:2502.18746* (2025).

- [6] Sarthak Dash, Gaetano Rossiello, Nandana Mihindukulasooriya, Sugato Bagchi, and Alfio Gliozzo. 2021. Open Knowledge Graphs Canonicalization using Variational Autoencoders. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP 2021)*. 10379–10394.
- [7] Anastasia Dimou, Miel Vander Sande, Pieter Colpaert, Ruben Verborgh, Erik Mannens, and Rik Van de Walle. 2014. RML: A generic language for integrated RDF mappings of heterogeneous data. *LDOW 2014* 1184 (2014).
- [8] Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitan, Robert Osazuwa Ness, and Jonathan Larson. 2024. From Local to Global: A Graph RAG Approach to Query-Focused Summarization. *arXiv preprint arXiv:2404.16130* (2024).
- [9] Oren Etzioni, Michele Banko, Stephen Soderland, and Daniel S Weld. 2008. Open information extraction from the web. *Commun. ACM* 51, 12 (2008), 68–74.
- [10] Michael Glass, Gaetano Rossiello, Md Faisal Mahbub Chowdhury, Ankita Naik, Pengshan Cai, and Alfio Gliozzo. 2022. Re2G: Retrieve, Rerank, Generate. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 2701–2715.
- [11] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. *arXiv preprint arXiv:2501.12948* (2025).
- [12] Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian, and Yuyu Yang. 2023. EvoPrompt: Connecting LLMs with Evolutionary Algorithms Yields Powerful Prompt Optimizers. (2023).
- [13] Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia d’Amato, Gerard de Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, et al. 2021. Knowledge Graphs. *ACM Computing Surveys (CSUR)* 54, 4 (2021), 1–37.
- [14] Pere-Lluís Huguet Cabot and Roberto Navigli. 2021. REBEL: Relation Extraction By End-to-end Language generation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*. 2370–2381.
- [15] Martin Josifoski, Marija Sakota, Maxime Peyrard, and Robert West. 2023. Exploiting Asymmetry for Synthetic Training Data Generation: SynthIE and the Case of Information Extraction. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. 1555–1574.
- [16] Omar Khattab, Arnab Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhaman, Saiful Haq, Ashutosh Sharma, Thomas T. Joshi, Hanna Moazam, Heather Miller, Matei Zaharia, and Christopher Potts. 2023. DSPy: Compiling Declarative Language Model Calls into Self-Improving Pipelines. *arXiv preprint arXiv:2310.03714* (2023).
- [17] Hanieh Khorashadizadeh, Fatima Zahra Amara, Morteza Kamaladdini Ezzabady, Frédéric Ieng, Sanju Tiwari, Nandana Mihindukulasooriya, Jinghua Groppe, Soror Sahri, Farah Benamara, and Sven Groppe. 2024. Research Trends for the Interplay between Large Language Models and Knowledge Graphs. In *Proceedings of Workshops at the 50th International Conference on Very Large Data Bases, VLDB 2024, Guangzhou, China, August 26–30, 2024*. VLDB.org.
- [18] Hanieh Khorashadizadeh, Nandana Mihindukulasooriya, Sanju Tiwari, Jinghua Groppe, and Sven Groppe. 2023. Exploring In-Context Learning Capabilities of Foundation Models for Generating Knowledge Graphs from Text. In *Text2KG at ESWC 2025*, Vol. 3447. 132–153.
- [19] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems*, Vol. 33. Curran Associates, Inc., 9459–9474.
- [20] Wenwu Li, Xiangfeng Wang, Wenhao Li, and Bo Jin. 2025. A Survey of Automatic Prompt Engineering: An Optimization Perspective. *arXiv preprint arXiv:2502.11560* (2025).
- [21] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing. 55, 9, Article 195 (Jan. 2023), 35 pages.
- [22] Huan Ma, Changqing Zhang, Yatao Bian, Lemao Liu, Zhirui Zhang, Peilin Zhao, Shu Zhang, Huazhu Fu, Qinghua Hu, and Bingzhe Wu. 2023. Fairness-guided few-shot prompting for large language models. In *NeurIPS 2023*. 43136–43155.
- [23] Igor Melynyk, Pierre Dognin, and Payel Das. 2022. Knowledge Graph Generation From Text. In *Findings of the Association for Computational Linguistics: EMNLP 2022*. Abu Dhabi, United Arab Emirates, 1610–1622.
- [24] Nandana Mihindukulasooriya, Sanju Tiwari, Carlos F. Enguix, and Kusum Lata. 2023. Text2KGBench: A Benchmark for Ontology-Driven Knowledge Graph Generation from Text. In *22nd International Semantic Web Conference ISWC 2023*, Vol. 14266. 247–265.
- [25] Sewon Min, Xinxin Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the Role of Demonstrations: What Makes In-Context Learning Work? *EMNLP 2022* (2022).
- [26] Itai Mondshine, Tzuf Paz-Argaman, and Reut Tsarfaty. 2025. Beyond english: The impact of prompt translation strategies across languages and tasks in multilingual llms. *LoResMT 2025 at NAACL 2025* (2025).
- [27] Krista Opsahl-Ong, Michael J Ryan, Josh Purtell, David Broman, Christopher Potts, Matei Zaharia, and Omar Khattab. 2024. Optimizing Instructions and Demonstrations for Multi-Stage Language Model Programs. *EMNLP 2024* (2024), 9340–9366.
- [28] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems* 35 (2022), 27730–27744.
- [29] Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language Models as Knowledge Bases?. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Hong Kong, China, 2463–2473.
- [30] Reid Pryzant, Dan Iter, Jerry Li, Yin Tat Lee, Chenguang Zhu, and Michael Zeng. 2023. Automatic Prompt Optimization with “Gradient Descent” and Beam Search. *EMNLP 2023* (2023), 7957–7968.
- [31] Libo Qin, Qiguang Chen, Fuxuan Wei, Shijue Huang, and Wanxiang Che. 2023. Cross-lingual Prompting: Improving Zero-shot Chain-of-Thought Reasoning across Languages. In *EMNLP 2023*. 2695–2709.
- [32] Kiran Ramnath, Kang Zhou, Sheng Guan, Soumya Smruti Mishra, Xuan Qi, Zhengyuan Shen, Shuai Wang, Sangmin Woo, Sullam Jeoung, Yawei Wang, et al. 2025. A systematic survey of automatic prompt optimization techniques. *arXiv preprint arXiv:2502.16923* (2025).
- [33] Gaetano Rossiello, Md. Faisal Mahbub Chowdhury, Nandana Mihindukulasooriya, Owen Cornec, and Alfio Massimiliano Gliozzo. 2023. KnowGL: Knowledge Generation and Linking from Text. In *AAAI AAAI Press*, 16476–16478.
- [34] Gaetano Rossiello, Nandana Mihindukulasooriya, Ibrahim Abdelaziz, Mihaela Bornea, Alfio Gliozzo, Tahira Naseem, and Pavan Kapanipathi. 2021. Generative relation linking for question answering over knowledge bases. In *Proceedings of the 20th International Semantic Web Conference (ISWC 2021)*. 321–337.
- [35] Satya S Sahoo, Wolfgang Halb, Sebastian Hellmann, Kingsley Idehen, Ted Thibodeau Jr, Sören Auer, Juan Sequeda, and Ahmed Ezzat. 2009. A survey of current approaches for mapping of relational databases to RDF. *W3C RDB2RDF Incubator Group Report 1* (2009), 113–130.
- [36] Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems* 33 (2020), 3008–3021.
- [37] Hao Sun, Alihan Hüyük, and Mihaela van der Schaar. 2024. Query-Dependent Prompt Evaluation and Optimization with Offline Inverse RL. In *The Twelfth International Conference on Learning Representations*.
- [38] Mirac Suzgun and Adam Tauman Kalai. 2024. Meta-Prompting: Enhancing Language Models with Task-Agnostic Scaffolding. *arXiv:2401.12954* (2024).
- [39] Edlira Vakaj, Nandana Mihindukulasooriya, Sanju Tiwari, and Sergio J. Rodriguez-Méndez. 2025. 4th International Workshop on Natural Language Processing for Knowledge Graph Construction (WWW ’25). Association for Computing Machinery, 2545–2548.
- [40] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: A Free Collaborative Knowledge Base. *Commun. ACM* 57, 10 (2014), 78–85.
- [41] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. *Advances in neural information processing systems* 35 (2022), 24824–24837.
- [42] Yuxin Wen, Neel Jain, John Kirchenbauer, Micah Goldblum, Jonas Geiping, and Tom Goldstein. 2023. Hard Prompts Made Easy: Gradient-Based Discrete Optimization for Prompt Tuning and Discovery. *Advances in Neural Information Processing Systems* 36 (2023), 51008–51025.
- [43] Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. 2022. An Explanation of In-context Learning as Implicit Bayesian Inference. *ICLR 2022* (2022).
- [44] Mert Yuksekgonul, Federico Bianchi, Joseph Boen, Sheng Liu, Pan Lu, Zhi Huang, Carlos Guestrin, and James Zou. 2025. Optimizing generative AI by backpropagating language model feedback. *Nature* 639 (2025), 609–616.
- [45] Zhao Zhang, Fuzhen Zhuang, Meng Qu, Fen Lin, and Qing He. 2018. Knowledge Graph Embedding with Hierarchical Relation Structure. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 3198–3207.
- [46] Lingfeng Zhong, Jia Wu, Qian Li, Hao Peng, and Xindong Wu. 2023. A Comprehensive Survey on Automatic Knowledge Graph Construction. *Comput. Surveys* 56, 4 (2023), 1–62.
- [47] Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitit, Harris Chan, and Jimmy Ba. 2023. Large Language Models are Human-Level Prompt Engineers. In *ICLR 2023*.
- [48] Kaijie Zhu, Jindong Wang, Jiaheng Zhou, Zichen Wang, Hao Chen, Yidong Wang, Linyi Yang, Wei Ye, Yue Zhang, Neil Gong, and Xing Xie. 2023. PromptRobust: Towards Evaluating the Robustness of Large Language Models on Adversarial Prompts. In *Proceedings of the 1st ACM workshop on large AI systems and models with privacy and safety analysis*. 57–68.