

CP-Model-Zoo: A Natural Language Query System for Constraint Programming Models (extended abstract)

Augustin Crespin ✉ 
UCLouvain, Belgium

Ioannis Kostis ✉ 
ICTEAM, UCLouvain, Belgium

Hélène Verhaeghe ✉ 
ICTEAM, UCLouvain, Belgium

Pierre Schaus ✉ 
ICTEAM, UCLouvain, Belgium

Constraint Programming (CP) and its high-level modeling languages have long been recognized for their potential to achieve the holy grail of problem-solving [5]. However, finding the right model is not an easy task.

First, the learning complexity of pure CP modeling languages (like MiniZinc [12], XCSP3 [2], or Essence [6]) can pose a challenge even for computer scientists. Second, even when the API is suggested in a mainstream language like CPMPy [8], modeling in CP is considered an art with many good practices or intuitions that are difficult to formalize. As explained in [13], one can employ techniques such as breaking symmetries, utilizing redundant constraints, and using multiple viewpoints with channeling constraints, among others. Finally, it is also known among CP experts that using global constraints (e.g., `element` [9], `table` [11],...) offers usually a stronger and more efficient pruning than the equivalent model with binary variables, as in MIP or SAT. Choosing the proper global constraint is also non-trivial.

With the recent progress in AI and large language models, we can envision a tool that generates expert-level CP models from a natural language description of a problem, thereby targeting a CP system. The reality is that we are not yet there and may not be ready to achieve this soon. When asked to model a combinatorial problem in CP, Large Language Models (LLMs) often suffer from the same usual recurring errors:

- syntactical errors: LLMs often make mistakes when directly asked to model in a less-known language (i.e., less-known compared to known worldwide, such as Java or Python)
- missing constraints or faulty models: LLMs generate text following the trained statistical distribution, without real reasoning, leading to such reasoning errors
- under-use of global constraints, as well as other known techniques such as redundant constraints or symmetry-breaking constraints

As CP experts, fixing such faulty models can sometimes take us more time and energy than starting from scratch. Learning quickly to reach a sufficient level of expertise in CP is, in our opinion, the safest approach to write good models. A CS professor arrived at a similar conclusion after experimenting with a programming task solved by students using LLM and without it [3]. One effective way to learn is by studying existing models that solve problems similar to the one at hand.

We propose a pragmatic approach that relies on the wealth of high-quality models accumulated over the years by the community through solver competitions [1, 14], CSPLib [7], and system tutorials [10]. The proposed solution retrieves the closest source code model from a database of models based on a user’s natural language description of a combinatorial optimization problem. By leveraging recent advancements in LLM technology, this approach is designed to prevent hallucinations, ensuring that only expert-validated models are presented to the user. It also eliminates the need for human intervention in data labeling. The system, called CP-Model-Zoo¹, is both fast and incremental. Adding a new model only requires generating a natural language description using an LLM and then computing a language model embedding in a vector space. Retrieving the most relevant model involves computing an embedding of the query and finding the model in the database with the closest similarity in the embedding space, as measured by cosine similarity.

Practically, the tool is based on a preprocessed database of CP models. For each of these models, several descriptions of the problem it solves are generated by an LLM. These descriptions are forced to be of various levels of expertise (we tested Novice, Intermediate, and Expert) to enhance the likelihood of matching with any user queries. Then, an embedding is generated from the code and the generated descriptions of the problem. These are the preprocessed data. Once a user makes a query, an embedding of it is produced and compared, using cosine similarity, to the preprocessed data, allowing the retrieval of the closest models, which are then proposed to the user.

To verify the usefulness of our platform, we perform some experiments. They rely on 67 Minizinc files from the MiniZinc examples², and the MiniZinc files present in CSPLib³ problems. The first experiment follows a leave-one-out approach between the query and each configuration for the embedding input, and tries to retrieve the model from which the description was generated. The second experiment uses CSPLib human-written descriptions and tries to retrieve the model it is associated with on the CSPLib website. As a metric for evaluation, we use Mean Reciprocal Rank (MRR) [4]. Through our experiments, we demonstrate that our system achieves excellent accuracy in retrieving the correct model based on user input descriptions simulated at different levels of expertise.

To facilitate broader access to the tool, a web-based application was developed as a demonstrative prototype. The application is hosted online⁴ and currently uses source code and Intermediate generated description as the embedding (i.e., the best-performing in our evaluation).

We believe that our tool, CP-Model-Zoo⁵, which requires minimal human intervention, offers interesting value for a broad audience. For CP learners, it can be used as an intelligent tutoring system to discover examples and learn how to model. For CP users, it can help retrieve known models for reuse. Furthermore, it could also serve as a starting point for an interactive LLM-aided modeling system, retrieving first the closest correct model and modifying it interactively with the user to address the exact complex problem at hand. The technical report details the complete architecture of CP-Model-Zoo in addition to the experiments done to assess its efficiency.

¹ Inspired by the Scheduling Zoo (<http://schedulingzoo.lip6.fr/>)

² <https://github.com/MiniZinc/minizinc-examples/tree/master>

³ <https://github.com/csplib/csplib>

⁴ <https://cp-model-zoo.info.ucl.ac.be/>

⁵ Source code: <https://github.com/crespina/LLM4CP>

References

- 1 Gilles Audemard, Christophe Lecoutre, and Emmanuel Lonca. Proceedings of the 2024 xcsp3 competition. *arXiv preprint arXiv:2412.00117*, 2024.
- 2 Frédéric Boussemart, Christophe Lecoutre, Gilles Audemard, and Cédric Piette. Xcsp3: an integrated format for benchmarking combinatorial constrained problems. *arXiv preprint arXiv:1611.03398*, 2016.
- 3 ACM Communications. The impact of ai on computer science education, 2023. Accessed: 2025-04-03. URL: <https://cacm.acm.org/news/the-impact-of-ai-on-computer-science-education/>.
- 4 Nick Craswell. *Mean Reciprocal Rank*, pages 1703–1703. Springer US, Boston, MA, 2009. doi:10.1007/978-0-387-39940-9_488.
- 5 Eugene C Freuder. In pursuit of the holy grail. *Constraints*, 2(1):57–61, 1997.
- 6 Alan M Frisch, Warwick Harvey, Chris Jefferson, Bernadette Martínez-Hernández, and Ian Miguel. E ssence: A constraint language for specifying combinatorial problems. *Constraints*, 13:268–306, 2008.
- 7 Ian P Gent and Toby Walsh. Csplib: a benchmark library for constraints. In *Principles and Practice of Constraint Programming–CP’99: 5th International Conference, CP’99, Alexandria, VA, USA, October 11–14, 1999. Proceedings 5*, pages 480–481. Springer, 1999.
- 8 Tias Guns. Increasing modeling language convenience with a universal n-dimensional array, cppy as python-embedded example. In *Proceedings of the 18th workshop on Constraint Modelling and Reformulation at CP (Modref 2019)*, volume 19, 2019.
- 9 Pascal Van Hentenryck and Jean-Philippe Carillon. Generality versus specificity: An experience with ai and or techniques. In *Proceedings of the Seventh AAAI National Conference on Artificial Intelligence*, pages 660–664, 1988.
- 10 Hakan Kjellerstrand. hakank/hakank, 2023. Accessed: 2025-04-03. URL: <https://github.com/hakank>.
- 11 Christophe Lecoutre. Str2: optimized simple tabular reduction for table constraints. *Constraints*, 16:341–371, 2011.
- 12 Nicholas Nethercote, Peter J. Stuckey, Ralph Becket, Sebastian Brand, Gregory J. Duck, and Guido Tack. Minizinc: Towards a standard cp modelling language. In Christian Bessière, editor, *Principles and Practice of Constraint Programming – CP 2007*, pages 529–543, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.
- 13 Barbara M. Smith. Modelling. In Francesca Rossi, Peter van Beek, and Toby Walsh, editors, *Handbook of Constraint Programming*, volume 2 of *Foundations of Artificial Intelligence*, pages 377–406. Elsevier, 2006. doi:10.1016/S1574-6526(06)80015-5.
- 14 Peter J Stuckey, Thibaut Feydy, Andreas Schutt, Guido Tack, and Julien Fischer. The minizinc challenge 2008–2013. *AI Magazine*, 35(2):55–60, 2014.

CP-Model-Zoo: A Natural Language Query System for Constraint Programming Models (technical report)

Augustin Crespin ✉ 
UCLouvain, Belgium

Ioannis Kostis ✉ 
ICTEAM, UCLouvain, Belgium

Hélène Verhaeghe ✉ 
ICTEAM, UCLouvain, Belgium

Pierre Schaus ✉ 
ICTEAM, UCLouvain, Belgium

Abstract

Constraint Programming and its high-level modeling languages have long been recognized for their potential to achieve the holy grail of problem-solving. However, the complexity of modeling languages, the large number of global constraints, and the art of creating good models have often hindered non-experts from choosing CP to solve their combinatorial problems. While generating an expert-level model from a natural-language description of a problem would be the dream, we are not yet there. We propose a tutoring system called *CP-Model-Zoo*, exploiting expert-written models accumulated through the years. CP-Model-Zoo retrieves the closest source code model from a database based on a user's natural language description of a combinatorial problem. It ensures that expert-validated models are presented to the user while eliminating the need for human data labeling. Our experiments show excellent accuracy in retrieving the correct model based on a user-input description of a problem simulated with different levels of expertise.

2012 ACM Subject Classification Computing methodologies → Natural language generation; Theory of computation → Constraint and logic programming; Information systems → Recommender systems

Keywords and phrases Constraint Identification, Large Language Models, Constraint Programming, Natural Language Processing, Information Retrieval, NLP4CP

Supplementary Material *Software (Source Code)*: <https://github.com/crespina/LLM4CP>
archived at [SoftwareHeritageIdentifier](#)

1 Introduction

Constraint Programming (CP) and its high-level modeling languages have long been recognized for their potential to achieve the holy grail of problem-solving: "*The user states the problem, the computer solves it*" [10]. However, finding the right model is not easy.

First, the learning complexity of pure CP modeling languages (like MiniZinc [27], XCSP3 [4], or Essence [11]) can pose a challenge even for computer scientists. Second, even when the API is suggested in a mainstream language like CPMPy [14], modeling in CP is considered an art with many good practices or intuitions that are difficult to formalize. As explained in [31], one can use techniques such as breaking symmetries, redundant constraint, using

multiple viewpoints with channeling constraints, etc. It is well known among CP experts that modeling with binary variables, as in MIP or SAT, generally leads to inefficient models for CP. Instead, efficient CP models usually use global constraints (e.g., `element` [17], `table` [23],...), offering a stronger and more efficient pruning. Choosing the proper one and their efficiency among many of them¹ is also non-trivial.

With the recent progress in AI and large language models, we can dream of a tool that generates expert-level CP models from a natural language description of a problem to target a CP system. The reality is that we are not yet there and may not be ready to achieve this soon for entirely new problems. As an example, we have tried feeding this toy puzzle description to state-of-the-art large language models (LLM): ChatGPT-4o, Gemini 2.5 Pro, and Mistral Lechat Pro:

You are given 13 words: BUOY, CAVE, CELT, FLUB, FORK, HEMP, JUDY, JUNK, LIMN, QUIP, SWAG, VISA, and WISH. These words collectively contain 24 unique letters. The goal is to assign each of these 24 letters to one of the six faces of four different cubes such that each letter appears on exactly one face of one cube. The assignment must ensure that for each word, its four letters are distributed across the four different cubes. Write a MiniZinc model for solving this problem.

```
include "alldifferent.mzn";

% The 24 unique letters found in the given words
set of int: LETTERS = 1..24;

% The 13 words given
array[1..13, 1..4] of int: WORDS = [
  | 1, 2, 3, 4 % BUOY
  | 5, 6, 7, 8 % CAVE
  | 9, 10, 11, 12 % CELT
  | 13, 14, 15, 16 % FLUB
  | 17, 18, 19, 20 % FORK
  | 21, 22, 23, 24 % HEMP
  | 3, 6, 12, 14 % JUDY
  | 10, 13, 19, 21 % JUNK
  | 4, 7, 9, 11 % LIMN
  | 8, 15, 16, 22 % QUIP
  | 5, 17, 20, 23 % SWAG
  | 1, 18, 24, 11 % VISA
  | 2, 9, 12, 16 % WISH
];

% Assign each letter to one of the four cubes
array[LETTERS] of 1..4: cube;

constraint alldifferent([cube[i] | i in LETTERS]);

% Ensure that each word's letters are on different cubes
constraint forall(w in 1..13) {
  cube[WORDS[w,1]] != cube[WORDS[w,2]] /\
  cube[WORDS[w,1]] != cube[WORDS[w,3]] /\
  cube[WORDS[w,1]] != cube[WORDS[w,4]] /\
  cube[WORDS[w,2]] != cube[WORDS[w,3]] /\
  cube[WORDS[w,2]] != cube[WORDS[w,4]] /\
  cube[WORDS[w,3]] != cube[WORDS[w,4]]
};
```

■ **Figure 1** Code produced by ChatGPT for the cube problem

The generated output for ChatGPT is given in Fig. 1 and is representative of the ones generated by concurrent services. This code has some serious issues:

- There are syntactical errors. It also didn't seem to understand the concept of decision variables and introduced plain sets instead.
- It introduced inconsistent data. Different integers were used to encode the same letters.
- It correctly tried to assign the letters of the same word to different cubes but used a decomposition instead of a global `alldifferent` constraint, which is not the best model an expert can imagine.
- The model does not enforce the cardinality of 6 faces per cube. This is a missing constraint. An expert would model it with a global cardinality constraint.

¹ The catalog of global constraints [3] count more than 400 constraints.

As CP experts, we tried to fix the model, but it took us more time and energy than starting from scratch, mainly because the bugs were not introduced incrementally. Learning quickly to reach a sufficient level of expertise in CP is, in our opinion, the safest approach to write good models. A CS professor arrived at a similar conclusion after experimenting with a programming task solved by students using LLM and without it [7]. One effective way to learn is by studying existing models that solve problems similar to the one at hand.

In this paper, we propose a pragmatic approach relying upon the wealth of high-quality models accumulated over the years by the community through solver competitions [32, 2], CSPLib [13], and system tutorials [21]. The proposed solution retrieves the closest source code model from a database of models based on a user’s natural language description of a combinatorial optimization problem. While utilizing recent advancements in LLM technology, this approach is designed to avoid hallucinations, ensuring that only expert-validated models are presented to the user. It also eliminates the need for human intervention in data labeling. The system, called CP-Model-Zoo², is both fast and incremental. Adding a new model only requires generating a natural language description using an LLM and then computing a language model embedding in a vector space. Retrieving the most relevant model involves computing an embedding of the query and finding the model in the database with the closest similarity in the embedding space, as measured by cosine similarity.

Through experiments, we demonstrate that our system achieves excellent accuracy in retrieving the correct model based on user input descriptions simulated at different levels of expertise. This paper details our methodology and introduces CP-Model-Zoo, an open-source tool available³ to the CP community and its potential users. We believe this tool offers a practical contribution to enhancing the learning and application of CP for a broader audience.

The paper is organized as follows. Section 2 presents the related work, providing an overview of existing research in the field. Section 3 outlines the methodology to rank and select the model starting from the natural language query. Section 4 conducts experiments to validate the tool’s capacity to retrieve relevant models from text descriptions of varying levels of expertise. Section 5 describes the tool’s user interface. Finally, Section 6 concludes the paper by summarizing the key findings and discussing future directions.

2 Related work

Different approaches have been tried to close the gap between CP formalism and natural language descriptions of CP problems.

First, the *NL4Opt* competition at *NeurIPS* in 2023 formulated the challenge of using *NLP* methods for extracting the meaning and formulation of an optimization problem based on its text description [30]. The top-performing approaches in the competition [12, 28, 19, 16] mainly relied on fine-tuned BERT [9] and BART [24] architectures.

Hahn et al. [15] fine-tuned the base version of the open-source language model T5 [29] to transform natural language descriptions into regex or other formal descriptions. Karia et al. [20] used LLMs to transform SAT formulas into natural language and vice-versa with a 95%

² Name inspired by the Scheduling Zoo (<http://schedulingzoo.lip6.fr/>), a tool to retrieve scientific articles based on the main characteristics of a scheduling problem.

³ <https://cp-model-zoo.info.ucl.ac.be/>

Tsouros et al. [34] tried to enhance the capacities of plain (not fine-tuned) LLMs with different techniques to make them CP modelers in CPMpy. The pipeline comprises a Named Entity Recognition to identify essential features of the problem, a relationship identifier between these named entities, and then a translation into CPMpy code. The system automatically corrects the errors in a loop, allowing users to refine the model iteratively. For modeling, a few LLM tricks are used, such as system prompting, chain or tree of thoughts, and plan-and-solve. The proposed method demonstrates promising results on selected examples; however, its generalizability to real-world datasets remains unverified.

Hotz et al. [18] made an LLM find solutions for the most well-known CP problems by iterating with the PyChoco API⁴ to find a syntactically correct solution and then a semantically correct one. The results were mitigated. Solutions were found for well-known problems (N-Queens, Map Coloring, TSP), but the system often hallucinated PyChoco functions, and the chats were unstable. Szeider used a similar idea in his MCP-Solver [33]. The client enables an LLM to interact with constraint models through natural language, while the server manages these interactions and translates them into MiniZinc operations using its Python API. A system prompt guides the LLM’s interaction. Likewise, solutions were found for famous problems, but no comparison with a plain LLM as a baseline was provided.

Michailidis et al. [26] used LLMs as CP solvers and modelers. Without surprises, using plain LLMs as solvers only yielded an accuracy of 11.46% on the *NL4Opt* dataset, showing clearly the need for a more specialized system to handle this task. For modeling, two techniques were used, i.e., blueprint model generation and in-context learning, with dynamically chosen examples. They were able to attain an 87.54 % solution accuracy on the *NL4Opt* dataset and 76.00% accuracy on a Logic Grid Puzzles dataset.

Let’s note that Voboril et al. [35] proposed StreamLLM, a system capable of generating streamliners (i.e., constraints added to a model to reduce the search space) in real time using LLMs. However, it relies heavily on the fact that a functional model already exists for the problem.

A more interactive approach, specifically a decision support for the meeting scheduling problem, was used by Lawless et al. [22]. An LLM is used in the Constraint Management component. It is prompted to select an action to take (add/delete a constraint, generate a suggestion, etc.). If a new constraint is added, two LLMs are used: an information checker and a coder in Python. It achieved around 80% accuracy for the information checking task and generated code that compiles correctly over 95% of the time and has high precision and recall (over 90%).

Similarly, Chen et al. [6] presented OptiChat, a natural language dialogue system intended to assist users in interpreting optimization model formulation, diagnosing infeasibility, doing a sensitivity analysis, retrieving data and evaluating modifications. Users upload a Pyomo-based optimization model [5], and can then interact with the model through natural language queries, asking about feasibility, sensitivity, etc. The evaluation is conducted on a dataset of Pyomo-based models with a success rate of above 80% for a majority of query types, with particularly strong performance in diagnosing infeasibility (87.2%) and sensitivity analysis (94.4%), while significantly reducing response time compared to human experts.

Contrary to related work, we do not use LLMs to directly model or solve the problem.

⁴ <https://pychoco.readthedocs.io/en/latest/>

Instead, we propose a system that takes input from a user’s description of a combinatorial problem and outputs the closest CP model from a database of models. Therefore, our system is closer in philosophy to an Information Retrieval (IR) system.

3 Methodology and problem description

You are given one or more MiniZinc models that represent a single classical constraint programming problem. Your task is to read the code and explain what the problem is about using very simple language. If there are several models for the same problem, do not explain each one separately. Instead, focus on explaining the overall problem. Assume the reader does not have much background in programming or mathematics. In your answer please explain:

The name of the problem.

What the problem is about in everyday terms.

What the main variables are and what they mean, using plain language.

What the basic restrictions or rules of the problem are, explained simply.

What the goal of the problem is (for example, what you want to minimize or maximize).

In your answer, do not include any introductory phrases (such as 'Here is the explanation of the problem')

Here is the source code:

```
-----
{source_code}
-----
```

(a) Novice Level - D1

You are experienced in constraint programming and familiar with MiniZinc. You are provided with one or more MiniZinc models representing a classic constraint programming problem. Your task is to identify the problem and explain it in clear, intermediate-level language. Assume the reader has some technical background but is not an expert. If there are several models for the same problem, do not explain each one separately. Instead, focus on explaining the overall problem. In your answer please explain:

The name of the problem.

A concise description of what the problem is about.

An explanation of the main decision variables and what they represent.

A description of the key constraints in plain language (avoid heavy mathematical notation).

An explanation of the problem’s objective (what is being minimized or maximized).

In your answer, do not include any introductory phrases (such as 'Here is the explanation of the problem')

Here is the source code of the model(s):

```
-----
{source_code}
-----
```

(b) Intermediate Level - D2

You are an expert in high-level constraint modeling and solving discrete optimization problems. In particular, you know MiniZinc. You are provided with one or several MiniZinc models that represents a single classical problem in constraint programming. Your task is to identify what is the problem modeled and give a complete description of the problem to the user. If there are several models for the same problem, do not explain each one separately. Instead, focus on explaining the overall problem. This is the source code of the model(s):

```
-----
{source_code}
-----
```

In your answer please explain:

name: The name of the problem

description: A description of the problem in English

variables: A string containing the list of all the decision variables in mathematical notation, followed by an explanation of what they are in English

constraints: A string containing the list of all the constraints in mathematical notation, followed by an explanation of what they are in English

objective: The objective of the problem (minimize or maximize what value)

In your answer, do not include any introductory phrases (such as 'Here is the explanation of the problem')

(c) Expert Level - D3

■ **Figure 2** Prompts to Generate Descriptions

Let $\mathcal{D} = \{s_1, s_2, \dots, s_N\}$ be a database of strings, where each s_i represents a source code implementation of a combinatorial optimization problem, possibly appended with an enriched textual description of the problem that is LLM-generated from the source code.

Given a query string q (a natural language problem description entered by the user), our goal is to retrieve the top- k most semantically similar strings from \mathcal{D} to the query string, therefore returning the most similar source-code implementation to help the user model their problem.

We use a pre-trained text embedding model to project strings into a shared embedding space. Let:

- $\phi : \mathcal{S} \rightarrow \mathbb{R}^d$ denote the embedding function, where \mathcal{S} is the string space and d is the embedding dimension.
- $\mathbf{e}_i = \phi(s_i)$ be the precomputed embedding vector of database string s_i .

- $\mathbf{e}_q = \phi(q)$ be the embedding vector of the query, computed at inference time.

Semantic similarity is measured via cosine distance in the embedding space:

$$\text{sim}(q, s_i) = \cos(\mathbf{e}_q, \mathbf{e}_i) = \frac{\mathbf{e}_q \cdot \mathbf{e}_i}{\|\mathbf{e}_q\| \|\mathbf{e}_i\|}$$

The top- k retrieved items are the $s_i \in \mathcal{D}$ with the highest similarity scores.

To construct \mathcal{D} , various strategies are possible:

1. **Code-only:** Raw source code strings.
2. **LLM-augmented:** Source Code + Synthetic Descriptions of varying expertise levels:

- **Novice:** Very simple explanations in non specific terms.

The prompt used to generate such descriptions can be found in Fig. 2a, while a sample *Novice*-level description for the knapsack problem is included in Fig. 3a.

- **Intermediate:** Domain-specific terminology.

Similarly, the prompt used to generate such descriptions can be found in Fig. 2b. In contrast, a sample *Intermediate*-level description for the knapsack problem is included in Fig. 3b.

- **Expert:** Formal mathematical formulations.

The prompt to be provided to an LLM to generate this level of descriptions can be found in Fig. 2c, while a sample *Expert*-level description for the knapsack problem is included in Fig. 4c.

At inference time, only \mathbf{e}_q requires computation, since the database embeddings $\{\mathbf{e}_i\}_{i=1}^N$ are precomputed. Retrieval reduces to:

$$\arg \max_{s_i \in \mathcal{D}}^k \text{sim}(\mathbf{e}_q, \mathbf{e}_i)$$

where $k \ll N$ is the number of models to return to the user.

Adding a new model to this framework is incremental and only requires 1) using an LLM to generate synthetic descriptions and 2) computing the corresponding embedding and storing it.

4 Experiments

Our experiments evaluate the capacity of our approach to rank highly relevant models. We designed an experimental setting that, starting from the set of models, does not require human input for the query while ensuring that the ground truth model related to the query is present.

All experiments were implemented⁵ using Python 3.10. **LlamaIndex** [25] was used for orchestration, while **Llama3-70b-8192** served as the LLM [1] to generate the synthetic descriptions. The LLM is hosted and served through an API provided by **Groq**⁶ to ensure

⁵ <https://github.com/crespina/LLM4CP>

⁶ <https://console.groq.com/>

```

**Knapsack Problem**

This problem is about packing objects of different sizes and values into a bag with a limited capacity. Imagine you're going on a trip and you have a small bag that can only hold so much weight. You have several items you want to bring, each with a different weight and importance to you. You want to pack the most important items that fit in the bag without exceeding its weight limit.

**Main variables:**

* 'n': the number of objects
* 'capacity': the maximum weight the bag can hold
* 'profit' (or 'value'): how important each object is to you
* 'size': how much each object weighs
* 'x': which objects to pack in the bag (how many of each object to take)

**Basic restrictions or rules:**

* You can't pack more objects than you have.
* The total weight of the packed objects can't exceed the bag's capacity.
* You can't pack a fraction of an object (you either take it or leave it).

**Goal:**

The goal is to pack the objects that give you the most importance (or value) while staying within the bag's weight limit.

```

(a) Novice Level

The problem is known as the 0/1 Knapsack Problem.

This problem is about selecting a subset of items to include in a knapsack of limited capacity, with the goal of maximizing the total value of the selected items. Each item has a size and a profit, and the knapsack has a limited capacity.

The main decision variables are the 'x' variables, which represent whether each item is included in the knapsack (1) or not (0).

The key constraints are:

- Each item can only be included once (or not at all), which is ensured by the 'x' variables being binary.
- The total size of the included items must not exceed the knapsack's capacity.

The objective is to maximize the total profit of the included items.

In other words, the problem is to find the optimal selection of items to put in the knapsack, given the capacity constraint, to maximize the total value of the selected items.

name: 0/1 Knapsack Problem

description: The 0/1 Knapsack Problem is a classic problem in combinatorial optimization. Given a set of objects, each with a weight and a value, determine the subset of objects to include in a knapsack of limited capacity to maximize the total value while not exceeding the knapsack's capacity.

variables: x (x[i] for i in OBJ or x[i] for i in 1..n)

The decision variables x represent the number of each object to include in the knapsack. In some models, x is a binary variable (0 or 1) indicating whether to include the object or not, while in others, x is an integer variable representing the quantity of each object to include.

constraints: x[i] >= 0, sum(i in OBJ)(size[i] * x[i]) <= capacity

The constraints ensure that the number of each object included is non-negative and that the total weight of the selected objects does not exceed the knapsack's capacity.

objective: maximize sum(i in OBJ)(profit[i] * x[i])

The objective is to maximize the total value of the objects included in the knapsack.

(b) Intermediate Level

(c) Expert Level

Figure 3 Description Samples for the Knapsack Problem

fast text generation. The embedding⁷ model is built upon the latest modernBERT pre-trained encoder-only foundation model [36], with an embedding vector size of $d = 768$ [37].

The experiments rely on two datasets origins: the **MiniZinc examples**⁸, and the **CSPLib**⁹ problems, which included both a MinZinc model and their corresponding human-written natural language descriptions. Among the 67 problems included in the dataset, 36 have a CSPLib description. We used only the Minizinc files as our database. The CSPLib descriptions were used as a validation set to verify whether our system could retrieve the correct Minizinc file (i.e., the associated Minizinc file on the CSPLib website was the truth value).

As a metric for evaluation, we use Mean Reciprocal Rank (MRR) [8], defined as: $MRR = \frac{1}{|Q|} \sum_{i=1}^{|Q|} \frac{1}{rank_i}$ where Q represent a set of queries and $rank_i$ refers to the rank position of the correct source-code model for the i -th query. An MRR of 1 means that the corresponding

⁷ <https://huggingface.co/Alibaba-NLP/gte-modernbert-base>

⁸ <https://github.com/MiniZinc/minizinc-examples/tree/master>

⁹ <https://github.com/csplib/csplib>

source-code model was ranked first for each query.

Different embeddings were created for each model, denoted as $SC + \{DX\}$. The first embedding, SC, is computed solely from the MiniZinc source code. The other embeddings $SC + \{DX\}$ are computed from the source code plus the LLM-generated text description with three different levels of expertise: Novice(1), Intermediate(2) and Expert(3). The prompts used to configure and generate them are displayed in Appendix 2.

The experiment follows a leave-one-out approach between the query and each configuration for the embedding input. The results in Table 1 report the MRR value computed on the ranking of the top-5 models retrieved out of the 67 models in total, which served as candidate models for selection, using different input of embeddings. These inputs included the source code alone, as well as the source code combined with LLM-generated descriptions of the problem at different levels of expertise. We tested four categories of queries. The first three correspond to the three levels of LLM-generated descriptions (67 queries for each level). They were compared to various embeddings containing a combination of the other descriptions (leave-one-out). The truth value here is the source code used to generate the description. The fourth one consists of the CSPLib human-written descriptions (36 queries), for which we have a truth value in the form of the Minizinc files associated with the website descriptions that were included in the source code database.

■ **Table 1** MRR values computed on the CP models, which served as candidate models for selection, using different sources of embeddings.

		Source of Embedding							
		SC	SC + D1	SC + D2	SC + D3	SC + D1&2	SC + D1&3	SC + D2&3	SC + D1&2&3
Query	D1	0.9851	-	1.0	0.9925	-	-	1.0	-
	D2	0.9254	0.9378	-	0.9577	-	0.9478	-	-
	D3	0.9527	0.9639	0.9925	-	0.9776	-	-	-
	CSPLib	0.8736	0.8898	0.9051	0.8676	0.8796	0.8704	0.8829	0.8634

The embeddings solely relying on MiniZinc source codes consistently perform worse than the others, supporting the hypothesis that attaching generated textual descriptions to the source code is beneficial. Notably, when *Novice*-level descriptions are used as queries, the embeddings containing both source codes and *Intermediate*-level descriptions, as well as the embeddings containing source code along with both *Intermediate*- and *Expert*-level descriptions, achieve an MRR of 1.0. This means that, for all problems, the corresponding model is correctly identified and ranked first every time. Furthermore, the scores remain consistently above 0.9, often exceeding 0.95. The quality of the results when the embeddings are based solely on the source code is noteworthy. While expecting lower MRR values in this configuration, we believe this can be partially attributed to some comments in the source code and the choice of variable identifiers, which add important semantic information that could aid the retrieval task.

Interestingly, despite expectations that embedding based on *Expert*-level descriptions would perform best when queried with CSPLib descriptions, this is not observed. Instead, the embeddings containing source code and *Intermediate*-level descriptions achieve the highest MRR of 0.9051. This is potentially due to CSPLib descriptions often being not truly high-level; for instance, they rarely incorporate extensive mathematical notation, whereas the generated *Expert*-level descriptions make heavy use of it. In other words, CSPLib descriptions

<p>Question: Dart competitions in pubs often have prizes for the first, second, third and fourth best throwers. What is the order of play for a darts tournament involving n throwers that:</p> <ul style="list-style-type: none"> - identifies the best k throwers (and their order); - has the least number of matches thrown; - has the least number of games on each of m dartboards (since tournaments typically run in parallel); - and is the most exciting (that is, prize winners should be discovered at the last possible minute so that everyone stays drinking beer until the end of the tournament). <p>Ranking:</p> <p>Rank 1: Problem Name: table_seating (Similarity: 0.6665750944113722)</p> <p>Rank 2: Problem Name: TTPPV (Similarity: 0.6545532118345632)</p> <p>Rank 3: Problem Name: golfers (Similarity: 0.6530544088881645)</p> <p>Rank 4: Problem Name: cluster (Similarity: 0.6466407509100225)</p> <p>Rank 5: Problem Name: doublechannel (Similarity: 0.6451558866521051)</p>	<p>Question: In the railway domain, the action of directing the traffic in accordance with an established timetable is managed by a software. However, in case of real time perturbations, the initial schedule may become infeasible or suboptimal. Subsequent decisions must then be taken manually by an operator in a very limited time in order to reschedule the traffic and reduce the consequence of the disturbances. They can for instance modify the departure time of a train or redirect it to another route. Unfortunately, this kind of hazardous decisions can have an unpredicted negative snowball effect on the delay of subsequent trains. Basically, the goal is to schedule adequately trains in order to bring them to their destination. The decision is to chose, for each train, which route it will follow and at what time.</p> <p>Ranking:</p> <p>Rank 1: Problem Name: bus_scheduling (Similarity: 0.670888769314417)</p> <p>Rank 2: Problem Name: doublechannel (Similarity: 0.649158749666409)</p> <p>Rank 3: Problem Name: constrained_connected (Similarity: 0.6463022454817552)</p> <p>Rank 4: Problem Name: nurses (Similarity: 0.6382525497004593)</p> <p>Rank 5: Problem Name: rcpsp (Similarity: 0.6369954887069154)</p>
---	---

(a) Darts tournament (CSPLib Prob 20)

(b) Train Traffic Rescheduling (CSPLib Prob 78)

<p>Question: Specification - Type 1</p> <p>Steel is produced by casting molten iron into slabs. A steel mill can produce a finite number, sigma, of slab sizes. An order has two properties, a colour corresponding to the route required through the steel mill and a weight. Given d input orders, the problem is to assign the orders to slabs, the number and size of which are also to be determined, such that the total weight of steel produced is minimised. This assignment is subject to two further constraints:</p> <p>Capacity constraints: The total weight of orders assigned to a slab cannot exceed the slab capacity.</p> <p>Colour constraints: Each slab can contain at most p of k total colours (p is usually 2).</p> <p>The colour constraints arise because it is expensive to cut up slabs in order to send them to different parts of the mill.</p> <p>The above description is a simplification of a real industrial problem. For example, the problem may also include inventory matching, where surplus stock can be used to fulfil some of the orders.</p> <p>Specification - Type 2</p> <p>The Type 1 specification does not constrain the number of slabs used. The Type 2 specification extends the objective to require further that the number of slabs used to accommodate the minimal weight of steel is also minimised.</p> <p>Ranking:</p> <p>Rank 1: Problem Name: carpet_cutting (Similarity: 0.6201213094705699)</p> <p>Rank 2: Problem Name: vessel-loading (Similarity: 0.6084444196269146)</p> <p>Rank 3: Problem Name: toomany (Similarity: 0.6037839776563689)</p> <p>Rank 4: Problem Name: car (Similarity: 0.6034512332968589)</p> <p>Rank 5: Problem Name: sonet_problem (Similarity: 0.6028041619673954)</p>	
--	--

(c) Steel Mill Slab Design (CSPLib Prob 38)

■ **Figure 4** Problems not in the database

tend to be long and verbose rather than direct. A similar issue arises with *Novice*-level descriptions, which include extensive explanations of fundamental concepts. One might also expect that the embeddings containing all description levels would perform best. However, a possible explanation for its under-performance is that incorporating *Expert*-level descriptions dilutes the overall representation by making it more verbose and redundant, reducing its conciseness.

We also entered CSPLib description queries for models not currently included in the CP-Model-Zoo database. Although this assessment is subjective, we found that the rankings produced by these queries were highly coherent and meaningful. The detailed results (using embedding SC+D2) of these queries are presented in the Fig. 4.

5 Web Application

To facilitate broader access to our research outcome by any interested party, a web-based application was developed as a demonstrative prototype, using the `gradio`¹⁰ Python framework.

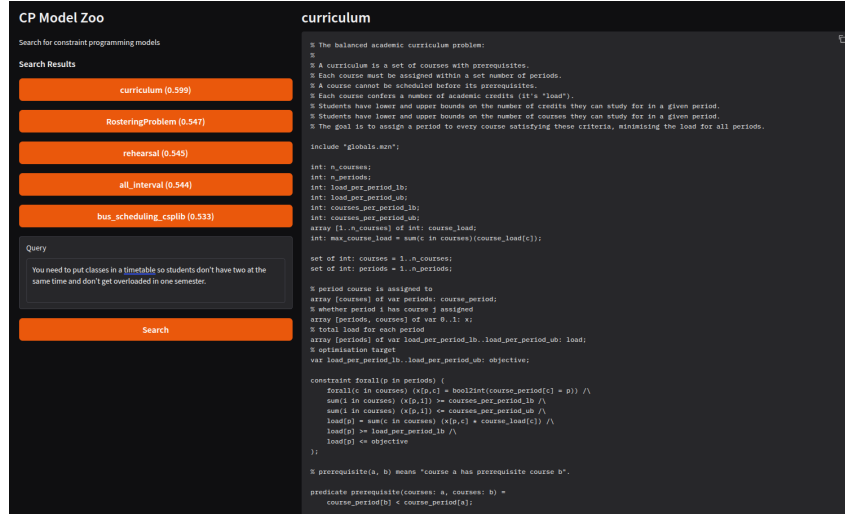
¹⁰<https://www.gradio.app/>

The application is hosted online¹¹ and uses as a knowledge base the SC+D2 embeddings set (see Section 4) for each combinatorial optimization problem available in it, 67 of these in total.

The end-user is able to interact with the tool by inputting a description of their problem in the textbox, found in the bottom left section of the application. Consecutively, the application retrieves and displays the ranked list of the top five most relevant CP models on the left-hand panel, above the textbox. Each model in this list is implemented as a clickable button, and upon selection, the corresponding MiniZinc source code is displayed on the right-side panel of the screen.

In Figure 5, we present a snapshot of the Graphical User Interface (GUI) of the application, upon execution of the following sample user query:

"You need to put classes in a timetable so students don't have two at the same time and don't get overloaded in one semester."



■ **Figure 5** Snapshot of the GUI of the web application, developed to facilitate access to the tool.

6 Conclusion & Future Work

In our work, we introduced CP-Model-Zoo, a tool designed to retrieve the closest source code model from a database based on a user's natural language description of a combinatorial optimization problem. Our system aims to connect problem descriptions made by a possibly non-CP expert with expert-validated CP models by using high-quality models accumulated over the years by the community.

Our experimental results demonstrate that CP-Model-Zoo achieves excellent accuracy in retrieving relevant models, even when the queries are simulated at different levels of expertise. We believe this tool, which requires minimal human intervention, offers interesting value as an intelligent tutoring system for CP. It could be hosted on the CSPLib, and the continuous addition of new clean models to it could be a community effort.

¹¹<https://cp-model-zoo.info.ucl.ac.be/>

References

- 1 AI@Meta. Llama 3 model card. 2024. URL: https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md.
- 2 Gilles Audemard, Christophe Lecoutre, and Emmanuel Lonca. Proceedings of the 2024 xcsp3 competition. *arXiv preprint arXiv:2412.00117*, 2024.
- 3 Nicolas Beldiceanu, Mats Carlsson, and Jean-Xavier Rampon. Global constraint catalog, 2010.
- 4 Frédéric Boussemart, Christophe Lecoutre, Gilles Audemard, and Cédric Piette. Xcsp3: an integrated format for benchmarking combinatorial constrained problems. *arXiv preprint arXiv:1611.03398*, 2016.
- 5 Michael L. Bynum, Gabriel A. Hackebeit, William E. Hart, Carl D. Laird, Bethany L. Nicholson, John D. Sirola, Jean-Paul Watson, and David L. Woodruff. *Pyomo — Optimization Modeling in Python*. Springer Cham, Springer Int Publishing Ag, Gewerbestrasse 11, Cham, Ch 6330, Switzerland., 2021.
- 6 Hao Chen, Gonzalo Esteban Constante-Flores, Krishna Sri Ipsit Mantri, Sai Madhukiran Kompalli, Akshdeep Singh Ahluwalia, and Can Li. OptiChat: Bridging optimization models and practitioners with large language models. URL: <http://arxiv.org/abs/2501.08406>, arXiv:2501.08406[cs], doi:10.48550/arXiv.2501.08406.
- 7 ACM Communications. The impact of ai on computer science education, 2023. Accessed: 2025-04-03. URL: <https://cacm.acm.org/news/the-impact-of-ai-on-computer-science-education/>.
- 8 Nick Craswell. *Mean Reciprocal Rank*, pages 1703–1703. Springer US, Boston, MA, 2009. doi:10.1007/978-0-387-39940-9_488.
- 9 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. URL: <http://arxiv.org/abs/1810.04805>, arXiv:1810.04805[cs], doi:10.48550/arXiv.1810.04805.
- 10 Eugene C Freuder. In pursuit of the holy grail. *Constraints*, 2(1):57–61, 1997.
- 11 Alan M Frisch, Warwick Harvey, Chris Jefferson, Bernadette Martínez-Hernández, and Ian Miguel. E ssence: A constraint language for specifying combinatorial problems. *Constraints*, 13:268–306, 2008.
- 12 Neeraj Gangwar and Nickvash Kani. Highlighting named entities in input for auto-formulation of optimization problems. In Catherine Dubois and Manfred Kerber, editors, *Intelligent Computer Mathematics*, pages 130–141. Springer Nature Switzerland. doi:10.1007/978-3-031-42753-4_9.
- 13 Ian P Gent and Toby Walsh. Csplib: a benchmark library for constraints. In *Principles and Practice of Constraint Programming–CP’99: 5th International Conference, CP’99, Alexandria, VA, USA, October 11-14, 1999. Proceedings 5*, pages 480–481. Springer, 1999.
- 14 Tias Guns. Increasing modeling language convenience with a universal n-dimensional array, cppy as python-embedded example. In *Proceedings of the 18th workshop on Constraint Modelling and Reformulation at CP (Modref 2019)*, volume 19, 2019.
- 15 Christopher Hahn, Frederik Schmitt, Julia J. Tillman, Niklas Metzger, Julian Siber, and Bernd Finkbeiner. Formal specifications from natural language. URL: <http://arxiv.org/abs/2206.01962>, arXiv:2206.01962[cs], doi:10.48550/arXiv.2206.01962.
- 16 JiangLong He, Mamatha N, Shiv Vignesh, Deepak Kumar, and Akshay Uppal. Linear programming word problems formulation using EnsembleCRF NER labeler and t5 text

- generator with data augmentations. URL: <http://arxiv.org/abs/2212.14657>, arXiv:2212.14657 [cs], doi:10.48550/arXiv.2212.14657.
- 17 Pascal Van Hentenryck and Jean-Philippe Carillon. Generality versus specificity: An experience with ai and or techniques. In *Proceedings of the Seventh AAAI National Conference on Artificial Intelligence*, pages 660–664, 1988.
 - 18 Lothar Hotz, Christian Bähnisch, Sebastian Lubos, Alexander Felfernig, Albert Haag, and Johannes Twiefel. Exploiting large language models for the automated generation of constraint satisfaction problems. URL: <https://www.semanticscholar.org/paper/Exploiting-Large-Language-Models-for-the-Automated-Hotz-B%C3%A4hnisch/f655f067a1d769d72ec48de52942f3cecf839db2>.
 - 19 Sanghwan Jang. Tag embedding and well-defined intermediate representation improve auto-formulation of problem description. URL: <http://arxiv.org/abs/2212.03575>, arXiv:2212.03575 [cs], doi:10.48550/arXiv.2212.03575.
 - 20 Rushang Karia, D. Dobhal, Daniel Bramblett, Pulkit Verma, and Siddharth Srivastava. Can LLMs translate SATisfactorily? assessing LLMs in generating and interpreting formal specifications. 2024. URL: <https://www.semanticscholar.org/paper/Can-LLMs-translate-SATisfactorily-Assessing-LLMs-in-Karia-Dobhal/e13f9819b30007181b3fd9d7d1dd76aa26e8b4df>.
 - 21 Hakan Kjellerstrand. hakank/hakank, 2023. Accessed: 2025-04-03. URL: <https://github.com/hakank>.
 - 22 Connor Lawless, Jakob Schoeffer, Lindy Le, Kael Rowan, Shilad Sen, Cristina St Hill, Jina Suh, and Bahareh Sarrafzadeh. "i want it that way": Enabling interactive decision support using large language models and constraint programming. version: 2. URL: <http://arxiv.org/abs/2312.06908>, arXiv:2312.06908 [cs], doi:10.48550/arXiv.2312.06908.
 - 23 Christophe Lecoutre. Str2: optimized simple tabular reduction for table constraints. *Constraints*, 16:341–371, 2011.
 - 24 Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880. Association for Computational Linguistics. URL: <https://aclanthology.org/2020.acl-main.703/>, doi:10.18653/v1/2020.acl-main.703.
 - 25 Jerry Liu. LlamaIndex, 11 2022. URL: https://github.com/jerryjliu/llama_index, doi:10.5281/zenodo.1234.
 - 26 Kostis Michailidis, Dimos Tsouros, and Tias Guns. Constraint modelling with LLMs using in-context learning. In Paul Shaw, editor, *30th International Conference on Principles and Practice of Constraint Programming (CP 2024)*, volume 307 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20:1–20:27. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISSN: 1868-8969. URL: <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.CP.2024.20>, doi:10.4230/LIPIcs.CP.2024.20.
 - 27 Nicholas Nethercote, Peter J. Stuckey, Ralph Becket, Sebastian Brand, Gregory J. Duck, and Guido Tack. Minizinc: Towards a standard cp modelling language. In Christian Bessière, editor, *Principles and Practice of Constraint Programming – CP 2007*, pages 529–543, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.

- 28 Yuting Ning, Jiayu Liu, Longhu Qin, Tong Xiao, Shangzi Xue, Zhenya Huang, Qi Liu, Enhong Chen, and Jinze Wu. A novel approach for auto-formulation of optimization problems. URL: <http://arxiv.org/abs/2302.04643>, arXiv:2302.04643[cs], doi:10.48550/arXiv.2302.04643.
- 29 Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. URL: <http://arxiv.org/abs/1910.10683>, arXiv:1910.10683[cs], doi:10.48550/arXiv.1910.10683.
- 30 Rindranirina Ramamonjison, Timothy T. Yu, Raymond Li, Haley Li, Giuseppe Carenini, Bissan Ghaddar, Shiqi He, Mahdi Mostajabdaveh, Amin Banitalebi-Dehkordi, Zirui Zhou, and Yong Zhang. NL4opt competition: Formulating optimization problems based on their natural language descriptions. URL: <http://arxiv.org/abs/2303.08233>, arXiv:2303.08233[cs], doi:10.48550/arXiv.2303.08233.
- 31 Barbara M. Smith. Modelling. In Francesca Rossi, Peter van Beek, and Toby Walsh, editors, *Handbook of Constraint Programming*, volume 2 of *Foundations of Artificial Intelligence*, pages 377–406. Elsevier, 2006. doi:10.1016/S1574-6526(06)80015-5.
- 32 Peter J Stuckey, Thibaut Feydy, Andreas Schutt, Guido Tack, and Julien Fischer. The minizinc challenge 2008–2013. *AI Magazine*, 35(2):55–60, 2014.
- 33 Stefan Szeider. MCP-solver: Integrating language models with constraint programming systems. URL: <http://arxiv.org/abs/2501.00539>, arXiv:2501.00539[cs], doi:10.48550/arXiv.2501.00539.
- 34 Dimos Tsouros, Hélène Verhaeghe, Serdar Kadioğlu, and Tias Guns. Holy grail 2.0: From natural language to constraint models. URL: <http://arxiv.org/abs/2308.01589>, arXiv:2308.01589[cs], doi:10.48550/arXiv.2308.01589.
- 35 Florentina Voboril, Vaidyanathan Peruvemba Ramaswamy, and Stefan Szeider. Realtime generation of streamliners with large language models. URL: <http://arxiv.org/abs/2408.10268>, arXiv:2408.10268[cs], doi:10.48550/arXiv.2408.10268.
- 36 Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Weller, Oskar Hallström, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, Nathan Cooper, Griffin Adams, Jeremy Howard, and Iacopo Poli. Smarter, better, faster, longer: A modern bidirectional encoder for fast, memory efficient, and long context finetuning and inference, 2024. URL: <https://arxiv.org/abs/2412.13663>, arXiv:2412.13663.
- 37 Xin Zhang, Yanzhao Zhang, Dingkun Long, Wen Xie, Ziqi Dai, Jialong Tang, Huan Lin, Bao-song Yang, Pengjun Xie, Fei Huang, et al. mgte: Generalized long-context text representation and reranking models for multilingual text retrieval. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 1393–1412, 2024.