

Generating Streamlining Constraints with Large Language Models

FLORENTINA VOBORIL, VAIDYANATHAN PERUVEMBA RAMASWAMY, and STEFAN SZEIDER, Technische Universität Wien, Austria

Streamlining constraints (or streamliners, for short) narrow the search space, enhancing the speed and feasibility of solving complex constraint satisfaction problems. Traditionally, streamliners were crafted manually or generated through systematically combined atomic constraints with high-effort offline testing. Our approach utilizes the generative capabilities of Large Language Models (LLMs) to propose effective streamliners for problems specified in the MiniZinc constraint programming language and integrates feedback to the LLM with quick empirical tests for validation. Evaluated across seven diverse constraint satisfaction problems, our method achieves substantial runtime reductions. We compare the results to obfuscated and disguised variants of the problem to see whether the results depend on LLM memorization. We also analyze whether longer offline runs improve the quality of streamliners and whether the LLM can propose good combinations of streamliners.

JAIR Track: Constraint Programming and Machine Learning

JAIR Associate Editor: Christian Bessiere

JAIR Reference Format:

Florentina Voboril, Vaidyanathan Peruvemba Ramaswamy, and Stefan Szeider. 2025. Generating Streamlining Constraints with Large Language Models. *Journal of Artificial Intelligence Research* 84, Article 16 (October 2025), 20 pages. DOI: [10.1613/jair.1.18965](https://doi.org/10.1613/jair.1.18965)

1 Introduction

Streamliners are certain constraints added to a constraint model to reduce the search space, thereby improving the feasibility and speed of finding solutions to complex constraint satisfaction problems. By incorporating domain-specific knowledge, streamliners can guide the constraint solver, allowing it to bypass less promising areas of the search space. Gomes and Sellmann (Gomes and Sellmann 2004) introduced streamliners to speed up the constrained-based search for hard combinatorial design problems. Today, streamliners are a standard tool for speeding up constrained-based search. Streamliners are closely related to implied/redundant constraints, symmetry-breaking constraints, and dominance-breaking constraints; however, adding a streamliner may even cause the constraint model to become inconsistent.

Originally, streamliners were hand-crafted by researchers who used their theoretical insight to analyze the constrained model. However, progress has also been made on the automated generation of streamliners (Spracklen, Dang, et al. 2023) by systematically trying the effect of some atomic constraints, such as imposing specific constraints on integer and function domains, like enforcing odd or even values, monotonicity, and properties like commutativity, as well as facilitating specific attributes in binary relations. These atomic restrictions are tested on thousands of problem instances, and those that show a good streamlining effect are systematically combined.

This paper proposes a different approach to automated streamliner generation that utilizes the generative capabilities (“creativity”) of Large Language Models (LLMs) to generate streamliners, providing a neuro-symbolic prototype implementation StreamLLM, and rigorously tests it on seven different constraint satisfaction problems

Authors’ Contact Information: Florentina Voboril, ORCID: <https://orcid.org/0009-0005-5683-5386>, fvoboril@ac.tuwien.ac.at; Vaidyanathan Peruvemba Ramaswamy, ORCID: <https://orcid.org/0000-0002-3101-2085>, vaidyanathan@ac.tuwien.ac.at; Stefan Szeider, ORCID: <https://orcid.org/0000-0001-8994-1656>, sz@ac.tuwien.ac.at, Algorithms and Complexity Group, Technische Universität Wien, Vienna, Austria.



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

© 2025 Copyright held by the owner/author(s).

DOI: [10.1613/jair.1.18965](https://doi.org/10.1613/jair.1.18965)

using at least 310 CPU days¹. Our approach leverages the capabilities of LLMs to infer potentially effective streamliners, similar to an experienced researcher’s intuitive grasp of a problem. By integrating LLMs into the streamliner generation process, we can potentially uncover unique and subtle patterns in the problem formulation and utilize them to speed up the solving process. To some extent, StreamLLM is closer related to a hand-crafted streamliner design by human experts rather than an automated bottom-up streamliner generation, the dominant strategy of the previous research (Spracklen, Dang, et al. 2023).

Our system StreamLLM for streamliner generation combines queries to LLMs with quick empirical validation on small test instances solvable within seconds. The system generates effective constraints within minutes through engineered prompts and adaptive feedback based on streamliner performance. This realtime approach allows streamliner benefits to be realized even for hard problem instances that would typically require hours to solve—any significant speedup can outweigh the brief time spent on streamliner generation.

We evaluate StreamLLM extensively across seven diverse constraint problems, including standard benchmarks and a novel hypergraph coloring problem, using hundreds of instances specified in MiniZinc. The problems are formulated in the MiniZinc constraint programming language (Nethercote et al. 2007). We test StreamLLM with two state-of-the-art LLMs (GPT-4o and Claude 3.5 Sonnet), different prompting variants, and different strategies for combining adaptive prompting with test runs on easy instances. Our results demonstrate remarkable effectiveness, with some streamliners reducing solving time by up to 99%.

To investigate whether these improvements stem from memorization or genuine problem understanding, we test StreamLLM on disguised and obfuscated variants of the problems. The system’s continued strong performance, particularly on problems with no known streamliners in the literature, suggests real analytical capabilities rather than mere pattern matching.

Beyond the realtime setting, we explore offline approaches with extended training periods and larger instance sets, achieving even better results in some cases. We also investigate StreamLLM’s ability to generate effective combinations of streamliners, finding that this often outperforms single constraints. Analysis of the generated streamliners reveals an interesting mix—some mirror expert-designed constraints while others take novel yet highly effective approaches.

Our experimental results suggest that LLM-generated streamliners can meaningfully reduce runtime across diverse constraint satisfaction problems with minimal problem-specific tuning, pointing to promising applications of AI assistance in constraint programming for computationally challenging problems. Part of this work appeared in preliminary and shortened form at the 1st International Workshop on Neuro-Symbolic Software Engineering (NSE) (Voboril, Ramaswamy, et al. 2025). *The source code and benchmark instances can be found on Zenodo (Voboril, P R, et al. 2025).*

2 Preliminaries

2.1 Constraint Programming

Constraint Programming (CP) is a methodology for solving combinatorial problems specified by means of declarative constraints. Please refer to Rossi, Van Beek, and Walsh (Rossi et al. 2006) for a comprehensive discourse. Examples include scheduling, routing, planning, etc. These problems can be of two types—*decision problems* requiring a yes/no answer, while *optimization problems* requiring you to find a solution that minimizes or maximizes a given objective function. In this paper, we focus only on decision problems.

MiniZinc (Nethercote et al. 2007) is a popular tool for solving CP problems. Problem specifications are written in the MiniZinc modeling language and are called *models*. MiniZinc compiles the model to a lower-level language and

¹This is a conservative estimate; we spent about 24 CPU days for instance generation, 75 CPU days for grading the instances, 162 CPU days for evaluating the performance of the generated streamliners, 20 CPU days for the offline approach, and 28 days for the combinations approach.

then solves it using one of many underlying solvers like Chuffed, Gecode, OR-Tools etc. Similarly, *Conjure* (Akgün, Frisch, Gent, et al. 2022; Akgün, Frisch, Hnich, et al. 2011) is also a tool for higher-level constraint programming, which uses *Essence* as its modeling language. In this paper, we use MiniZinc as the modeling language of choice, while *Conjure* is used in the instance generation pipeline AutoIG. More specifically, the structure and format of the instances are specified in *Essence* and given to *Conjure*, which then generates instances following these constraints.

2.2 Streamliners

Streamlining constraints or *streamliners*, introduced by Gomes and Sellmann (Gomes and Sellmann 2004), are constraints that are added to the constraint programming models with the goal of narrowing the focus to a small but highly promising segment of the search space. Thus, they have the potential to significantly reduce the search effort for hard combinatorial problems. For example, Gomes and Sellmann (Gomes and Sellmann 2004) used constraints for “Latin squareness” as a streamliner for finding Diagonally Ordered Magic Squares. Streamliners have been successfully applied in a diverse range of settings, such as combinatorial design (Díaz et al. 2017; LeBras et al. 2012; Smith et al. 2005), Ramsey-type problems and discrepancy (Bras et al. 2014; Liu et al. 2021), SAT/CSP solver design (Ansótegui et al. 2022; Gomes, Sabharwal, et al. 2006; Grover et al. 2018; Heule et al. 2019), and automatically generated streamliners (Spracklen, Dang, et al. 2023, 2020).

We note that streamlining constraints are not required to be sound. This means that adding the streamlining constraint may make a satisfiable instance of a constraint model unsatisfiable. As a consequence, an unsatisfiable instance in the streamlined model does not imply that the instance is unsatisfiable in the original model.

Other constraints that are similar to streamliners are *implied* (or *redundant*) constraints, *symmetry-breaking* constraints, and *dominance-breaking* constraints. Implied constraints do not change the set of feasible solutions (Charnley et al. 2006; Colton and Miguel 2001; Frisch, Jefferson, and Miguel 2004; Frisch, Miguel, et al. 2003, 2001). Symmetry-breaking constraints eliminate certain solutions within each equivalence class while ensuring that at least one solution from each class remains (Fichte et al. 2020; Flener et al. 2001; Frisch, Hnich, et al. 2009, 2006; Frisch, Jefferson, Martinez-Hernandez, et al. 2007; Itzhakov and Codish 2022). Dominance-breaking constraints are applicable in the context of optimization problems (possibly formulated as a decision problem with the objective value explicitly stated in the model), as they disallow solutions that are known to be suboptimal. It might disallow optimal solutions as well, as long as at least one optimal solution remains (Chu and Stuckey 2015; Gomes and Sellmann 2004; J. H. M. Lee and Zhong 2022; Prestwich and Beck 2004). In contrast to streamliners, such constraints are guaranteed to be sound.

Finding a useful streamliner manually by a human expert is a time-consuming process. Therefore, it is appealing to generate streamliners automatically, a direction explored successfully in previous work (Spracklen, Akgün, et al. 2018; Spracklen, Dang, et al. 2023, 2019; Wetter et al. 2015). A common point between all these previous approaches is that they treat the streamliner generation as a high-effort, offline task, taking up to 4 CPU days for a single problem. The streamliner is built up systematically from elementary constraints to reduce the domain of a decision variable and tested on a large set of automatically generated instances.

Similar to other recent studies (Spracklen, Dang, et al. 2023), we use the term streamliner in a wider sense to also accommodate redundant, symmetry-breaking constraints. This is particularly useful in our context since, for automatically constructed constraints, it is difficult to determine which type the new constraint is. At the same time, we acknowledge that the practical impact of different constraint types can vary substantially. Therefore, rather than focusing on a strict theoretical classification, we adopt a pragmatic viewpoint and evaluate each generated constraint empirically.

2.3 Large Language Models (LLMs)

Large Language Models (LLMs) are advanced AI systems based on the transformer models (Vaswani et al. 2017) and trained on vast data sets to produce human-like text and source code in response to instruction prompts (Minaee et al. 2024). These models, trained on vast amounts of text data, can produce human-like text across various domains and styles (Brown et al. 2020). Recent advancements have expanded their capabilities beyond traditional language tasks, including generating, analyzing, and debugging code across multiple programming languages (Chen et al. 2021; Pei et al. 2023; Wu et al. 2023; Xu et al. 2022). In addition to code-related tasks, LLMs have shown promise in mathematical reasoning and problem-solving (Lample et al. 2022; Romera-Paredes et al. 2024); they can process and generate mathematical expressions, solve equations, and even assist with proofs. However, it is essential to note that while LLMs can produce seemingly correct mathematical output, their responses should be carefully verified. The models' performance on these tasks varies, and they may sometimes generate plausible-looking but incorrect solutions (Polu et al. 2023). Despite these limitations, the potential applications of LLMs in fields such as computer science, mathematics, and engineering are substantial and continue to expand as the technology evolves.

3 LLM-based Streamliner Generation and Validation

In this section, we discuss the different approaches we devise to generate and test streamliners. The key difference between our approach and the past work by Spracklen et al. (Spracklen, Dang, et al. 2023) is that instead of deriving streamliners by composing and combining atomic constraints, we synthesize entire streamliners in one go. We try several variants of our approach, one which is entirely realtime and requires no precomputation (unlike Spracklen et al. (Spracklen, Dang, et al. 2023)), another which can directly synthesize combinations of streamliners, and one which is offline. We evaluate our approach on benchmark problems from the literature along with modified versions of those problems to give insights into the strengths and weaknesses of our approach.

3.1 Realtime Approach

The realtime approach utilizes the generative capabilities of LLMs (usually taking only a few seconds) to propose and validate streamliners in real time. Such on-the-fly generation is not possible with high-effort streamliner generation that builds a streamliner systematically from elementary steps and takes several CPU days (Spracklen, Dang, et al. 2023).

Our system submits several queries (prompts) to the LLM that result in candidate streamliners and tests the candidate streamliners on a small set of n_{train} easy *training instances* that can be solved in less than t_{train} seconds. Rigorous testing on several problems indicates that a relatively small number of test instances suffices (about 15) and even small and easy test instances (solvable in under 10 seconds with the unstreamlined model) provide a good indication of how well streamliners will work on large and hard instances; we lay out these experiments in more detail in Sections 4.5 and 4.6. We test various strategies for this type of streamliner generation, which we divide into static and adaptive categories. Algorithm 1 outlines the general strategy, which can be summarized as follows. First, the baseline solving times of the training instances are determined by solving them with the original, unstreamlined model. These baseline times are later used as reference timeouts. Then, for about ten minutes, the algorithm repeatedly prompts the LLM to generate candidate streamliners and evaluates their performance on the training instances. Streamliners that cause errors or result in slower solving times than the baseline are discarded, while the promising ones are kept. In the adaptive mode, the evaluation results are also appended as feedback to the prompt, allowing the LLM to refine its suggestions based on past performance. To maintain diversity, every third iteration is performed without feedback, leaving room for new ideas. After the ten minutes have passed, the algorithm selects the best-performing streamliners from all evaluated candidates and outputs them.

Algorithm 1 General Strategy

Input: an unstreamlined MiniZinc model \mathcal{M} , a set of training instances \mathcal{I} , a prompt \mathcal{P} , an LLM \mathcal{L} , integers n, k **Output:** a set of streamliners \mathcal{T}

- 1: $\mathcal{S} \leftarrow \emptyset, \mathcal{S}' \leftarrow \emptyset, \mathcal{P}' \leftarrow \mathcal{P}$
 - 2: Find the baseline solving times $t_{\text{base}}(I)$ for each instance $I \in \mathcal{I}$, by solving them with the unstreamlined model \mathcal{M} .
 - 3: **repeat**
 - 4: Obtain a set \mathcal{S}' of n streamliners by prompting the LLM \mathcal{L} with prompt \mathcal{P}' .
 - 5: Evaluate the performance of the streamliners from \mathcal{S}' in parallel on the training instances using the baseline times t_{base} as the corresponding timeouts.
 - 6: Discard from \mathcal{S}' the streamliners that produce errors or time out on all the training instances.
 - 7: $\mathcal{S} \leftarrow \mathcal{S} \cup \mathcal{S}'$
 - 8: Set $\mathcal{S}' \leftarrow \emptyset$ every third iteration. *// only relevant for Adaptive Mode*
 - 9: **if** Adaptive Mode **then**
 - 10: Construct the new prompt \mathcal{P}' by appending feedback about the streamliners in \mathcal{S}' to the prompt \mathcal{P} . This feedback includes information such as whether they produced errors or unsatisfiable instances; or if their performance was better or worse than the unstreamlined model.
 - 11: **end if**
 - 12: **until** 10 minutes have passed.
 - 13: Find a set \mathcal{T} of k streamliners from \mathcal{S} that collectively achieve the best performance over the training instances.
 - 14: **return** \mathcal{T}
-

Our realtime approach targets scenarios where a user needs to solve hard constraint problems that typically require hour-long solve times. Rather than directly attempting these hard instances, the user provides a few small instances to our system. This proposes efficient streamliners after interacting with an LLM and a constraint solver within minutes. The user can then run these streamlined versions of the original model in parallel with the unstreamlined version, potentially achieving significant speedup that outweighs the initial streamliner generation time.

3.2 Prompt Engineering

The prompt given in Figure 1, is the prompt we use for most of our experiments. In the beginning, the task is shortly summarized. Then, following the Chain of Thought technique (Wei et al. 2022), the task is split into single steps that are explicitly explained. In the end, there are some compliance rules to ensure the quality and the correct format of the response. Especially the correct JSON output format is important for the automated tests on the training instances. We decided on this prompt because it showed good results in our preliminary experimentation (Voboril, Ramaswamy, et al. 2025).

3.3 Obfuscation and Disguise

The capabilities of LLMs to generate original ideas versus relying on *memorization* remains a key debate in AI research (J. Lee et al. 2023; McCoy et al. 2023). This question is particularly relevant for streamliner generation, as the practical value would be limited if LLMs could only identify streamliners for well-documented problems where such techniques are already published and likely included in their training data. To investigate this capability empirically, we examine two problem variants: disguised and obfuscated versions.

Objective

Analyze the given MiniZinc code and suggest five additional constraints to enhance the problem-solving process. These constraints can include streamlining, implied, symmetry-breaking, or dominance-breaking constraints.

Steps

- (1) **Analyze Content:** Read the provided MiniZinc code. Understand the problem being addressed, including its variables, constraints, and optimization goals.
- (2) **Generate additional Constraints:** Based on your analysis, create five unique constraints. These should offer targeted modifications or restrictions designed to reduce the search space effectively.
- (3) **Always return your constraints as a JSON object, adhering to the structure:**

```
{“streamliner_1”: “constraint <MiniZinc constraint>”, ..., “streamliner_5”: “constraint <MiniZinc constraint>”}.
```

Your final output should exclusively be the JSON object containing the five constraints.
- (4) As a response, you will get feedback for each constraint. Some constraints might lead to errors, timeouts, or unsatisfiable instances.
- (5) Use the information provided in the previous step to generate five new, and hopefully better constraints.
- (6) Repeat steps 3 to 5 multiple times.

Compliance Rules

- (1) **Response Format:** Your final output should exclusively be the JSON object containing the five constraints, adhering to the structure:

```
{“streamliner_1”: “constraint <MiniZinc constraint>”, ..., “streamliner_5”: “constraint <MiniZinc constraint>”}.
```

Do not forget the semicolon at the end of the constraint!
- (2) **Code Quality:** All MiniZinc code provided for the constraints must be syntactically correct and functional. For some functions you may need to include an additional library.
- (3) **Creativity:** You’re encouraged to be innovative in proposing constraints, keeping in mind their purpose: to narrow down the search space efficiently without oversimplifying the problem.

Fig. 1. Base Prompt

To *disguise* a problem, we rename all occurrences of all identifiers such that they express the semantics of a different problem while maintaining the syntax and the structure of the original problem. We also insert comments describing the semantics of the new problem. For instance, to disguise the well-known Social Golfers problem as a newly invented “Zoo Habitat Rotation” problem, we map Golfer \mapsto Animal, Group \mapsto Habitat, and Week \mapsto Season. Additionally, we add the following comment at the top of the MiniZinc model:

“The zookeepers at a large zoo are organizing seasonal habitat rotations to encourage socialization among animals. Each season, animals are grouped into mixed-species habitats of a fixed size. The zookeepers want to ensure that no two animals share a habitat more than once during the rotation schedule, while maximizing the number of season the rotations can last. This problem seeks for an allocation of animals to habitats which ensures maximum variety in animal pairings across seasons while maintaining the integrity of habitat sizes.”

Obfuscating a problem is similar to disguising, except that instead of mapping to names with the semantics of a different problem, we map to arbitrary names like `id8`, `id12`, and `id5`, and we strip all comments from the model to avoid revealing the original problem. Thus, for each original problem in our benchmark, we introduce two new problems to our benchmark by disguising and obfuscating the model (`.mzn` file) and instances (`.dzn` files) of the original problem. We then generate and evaluate streamliners for these problems using the same prompts as the original problem.

3.4 Combinations of Constraints

Previous work by Spracklen et al. (Spracklen, Dang, et al. 2023) demonstrates that combining multiple streamlining constraints can be more effective than using them individually. To evaluate how combinations of streamliners perform in our StreamLLM approach, we provide another prompt alternative to the base prompt. It was created by modifying the LLM base prompt to output combinations of constraints rather than single constraints, with corresponding adjustments to the required JSON format. We evaluate the results with this alternative prompt in Section 4.9. To establish a baseline for our approach, we evaluate Spracklen et al.’s streamliner combinations on our benchmark instances, with detailed results presented in Subsection 4.10.

3.5 Offline Approach

The offline approach investigates whether extended training time improves streamliner quality. We increase the timeout from 10 minutes to 4 hours and expand the training set from 15 instances solvable within 10 seconds to 35 instances with solving times between 1 and 30 seconds. The process follows the static approach outlined in Algorithm 1. In each iteration, we randomly pick GPT-4o or Claude 3.5 Sonnet, using either the original or combination-generating prompt described in Section 3.4 to leverage the strengths of different LLMs and prompting strategies.

4 Experiments

All instances, MiniZinc models, and the Python implementation of StreamLLM (which includes the prompts) are available on Zenodo².

4.1 Setup and Hardware

We use MiniZinc 2.8.3 and the Chuffed 0.13.1 solver. As LLMs we use GPT-4o (`gpt-4o-2024-11-20`) and Claude 3.5 Sonnet (`claude-3-5-sonnet-20241022`) and access them via the `openai 1.29.0` and `anthropic 0.25.8` packages in Python 3.11.5. We evaluate the running times of the test instances on compute nodes with two 2.4 GHz 10-core Intel Xeon E5-2640 v4 processors with 80 GB of RAM each.

4.2 Benchmark Problems

Although our method can be applied to optimization problems, similar to previous work by Spracklen et al. (Spracklen, Dang, et al. 2023), we focus here on decision problems. This way we can evaluate streamliners only on the basis of their running times. For optimization problems, one needs to consider solution quality

²<https://zenodo.org/doi/10.5281/zenodo.13331048>

as well. We construct a benchmark set of seven constraint satisfaction problems to test the LLM-generated streamliners. This includes four problems which are already considered by Spracklen et al. (Spracklen, Dang, et al. 2023), namely **Balanced Incomplete Block Design (BIBD)**^{†‡}, **Car Sequencing (CS)**^{†§}, **Social Golfers (SG)**^{†‡}, and **Vessel Loading (VL)**^{†‡}. In addition to that, we include two other well-known problems, **Black Hole (BH)**^{†§} and **Carpet Cutting (CC)**^{*§}. The main criterion for selecting these six problems is that the MiniZinc model must be readily available. Further filtering is done based on the ease of generation of instances of the desired difficulty. Note that, if present, we strip any redundant constraints, symmetry-breaking constraints, and excessive documentation from the MiniZinc models of these problems. This is the case for CC and SG. The only exception is BIBD, for which, without symmetry-breaking, we are unable to find problem instances of the desired difficulty.

Finally, we added the well-known NP-hard **Hypergraph Coloring (HC)** problem (Garey and Johnson 1979; Lovász 1973) (described below), for which, to the best of our knowledge, no constraint model is available. The main reason for including this problem is to see whether our approach relies on memorization (e.g., Lee et al. (J. Lee et al. 2023) and McCoy et al. (McCoy et al. 2023)).

Hypergraph Coloring (HC): Given integers c and m , a set V of vertices, and a set E of hyperedges, where each hyperedge is a subset of V , find a coloring of vertices using at most c colors such that no hyperedge is monochromatic, i.e., is only incident to vertices of the same color. Further, the *imbalance*, i.e., the difference in sizes of the largest and the smallest color class, must not exceed m .

4.3 Instance Generation

Since we work with a broad range of problems, working with instances of similar/uniform difficulty ensures the generality of our findings and prevents bias from creeping in due to any one particular problem. Hence we construct our own data set by generating instances of desired difficulties. For each problem, we create 15 *training instances* that take less than 10 seconds. Further, we create at least 50 *test instances* that take between 10 minutes and 2 hours to solve. The 10-minute lower limit is because we want our approach to be comparable in a realtime setting. The 2-hour upper limit is due to limited computational resources.

We generated satisfiable instances of the desired difficulties for all the above problems. Since SG, BH, and BIBD all have relatively simple input specifications, i.e., a small number of integer parameters, we generated instances exhaustively and tested their difficulty. For CC, CS, HC, VL we generated instances using the AutoIG pipeline developed by Dang et al. (Dang et al. 2022). This involved coming up with a parameterized Essence specification describing the input space of each problem. We then fix a difficulty window in the form of lower and upper bounds for the MiniZinc running time and then use the automatic algorithm configuration tool *irace* to find an optimal configuration of those parameters such that the constraint-solving tool *Conjure* can find an instance that is very likely to fall in the desired difficulty window.

A detailed visualization of the distribution of the original running times (or in other words, difficulty levels) of the generated instances is shown in Figure 2. The total number of instances for each problem is stated in the legend. Overall, the compiled CSP instances range from tens of variables and constraints to millions of variables and constraints with an average of around 200,000 variables and around 300,000 constraints. We list the ranges of the problem-specific parameters of our benchmark instances below.

[†]Problem description and references available at <https://www.csplib.org/Problems>

[‡]Model from <https://www.csplib.org/Problems>

[§]Model from <https://github.com/MiniZinc/minizinc-benchmarks>

^{*}Problem described by Schutt et al. (Schutt et al. 2011)

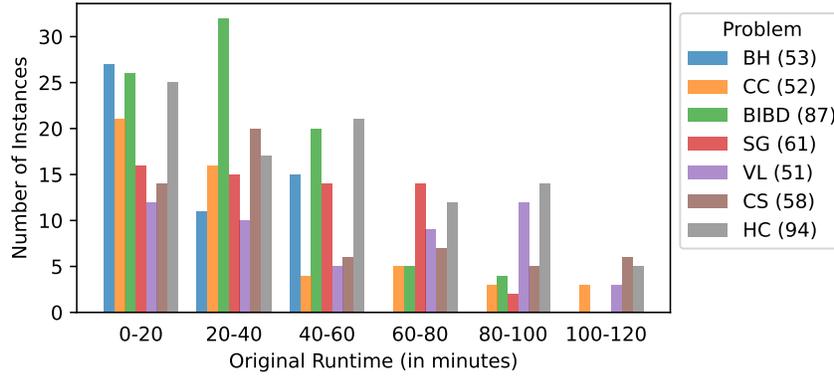


Fig. 2. Histogram showing the number of instances for each problem, sorted by their original running times and partitioned into 20-minute intervals. Further, the total number of instances for each problem is shown next to the problem name in the legend.

Black Hole: Each instance is just a different way of splitting and permuting a deck of 52 cards into 3 stacks.

Carpet Cutting: Number of room $\in [1, 6]$, number of room rectangles $\in [4, 38]$, number of stairs $\in [1, 5]$.

Balanced Incomplete Block Design: $v \in [9, 46]$, $k \in [2, 13]$, $\lambda \in [1, 15]$.

Social Golfers: Number of groups $\in [6, 38]$, number of golfers $\in [21, 174]$, number of weeks $\in [2, 27]$.

Vessel Loading: Number of classes $\in [1, 5]$, number of containers $\in [19, 67]$, vessel area $\in [153, 512]$.

Car Sequencing: Number of cars $\in [20, 48]$, number of configurations $\in [13, 20]$, number of options $\in [3, 6]$.

Hypergraph Coloring: $|V| \in [20, 100]$, $|E| \in [10, 99]$, number of colors $\in [3, 5]$.

4.4 k-Best Selection

As already explained in Section 2.2, there is no guarantee that every satisfiable instance in the original model is also satisfiable in a streamlined model. It also might happen that one streamliner works well on the training instances but is impractical for the large test instances. In order to get more robustness, we decided to not only rely on one streamliner but also to return k streamliners that work best on the training instances. In the experiments, we run the original model and $k = 3$ streamlined models in parallel and stop as soon as any of the $k + 1$ models have found a solution.

4.5 Preliminary Experiment A

The objective of this experiment is to determine a suitable maximal running time t_{train} of the unstreamlined model for the training instances. We want to keep t_{train} as low as possible so that the evaluation can be done quickly but still provide good predictions for the streamliner performance on the significantly larger test instances.

We run the experiment with the problems BIBD, BH, and SG. For each problem, we let the LLM suggest ten streamliners. Then, we pick those three streamliners that perform best on the training instances, and among the three, we pick the one that performs best on the test instances. We use five sets of training instances with an upper bound of $t_{\text{train}} \in \{10, 20, 30, 60\}$ seconds, respectively. The experiments show that t_{train} has no influence on the one streamliner that was picked in the end, hence we settled on the upper bound $t_{\text{train}} = 10$ for the following experiments. To reduce the influence of I/O operations, we require training instances to take at least 1 second to be solved.

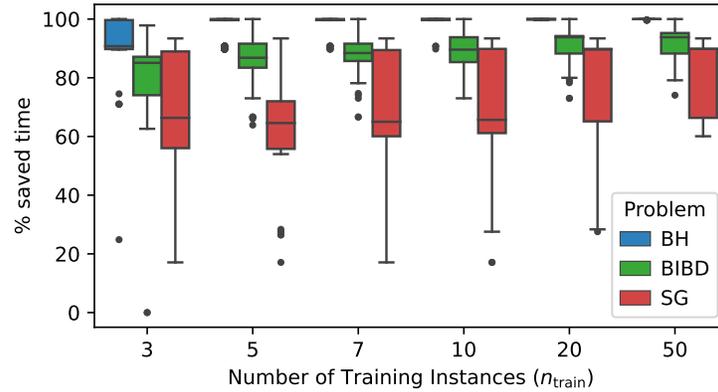


Fig. 3. Normalized saved times for different number of training instances for three problems

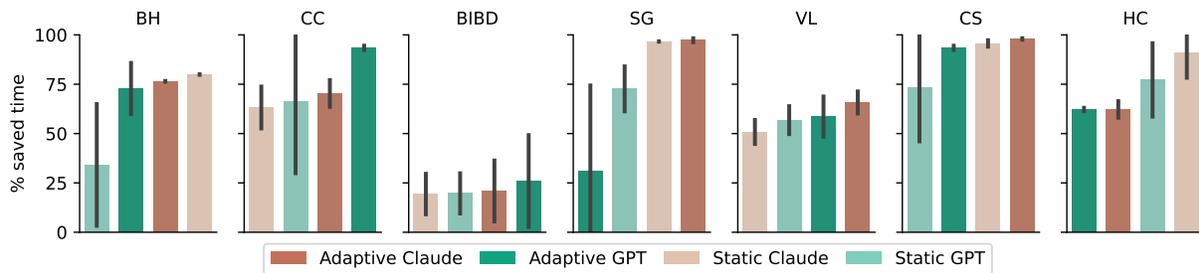


Fig. 4. Percentage of reduction in solving time with both static and adaptive approaches using LLMs Claude and GPT across seven problems. Each bar and black line denotes, respectively, the mean and standard deviation of two runs.

4.6 Preliminary Experiment B

The goal of this experiment is to decide on the number n_{train} of training instances. More training instances promise better results on the test instances but make the evaluation process longer, which we want to avoid for realtime streamliner generation. Hence we aim at a fair compromise.

We conduct the experiment with the same three problems as in Experiment 1a. For each problem, we generate ten streamliners. For each $n_{\text{train}} \in \{3, 5, 7, 10, 20, 50\}$, we randomly pick n_{train} training instances. We determine the combination of three out of the ten streamliners that perform best on the training instances and compute the time this combination saves on the test instances, normalized by the time saved by the virtually best combination of three streamliners. We repeat this 100 times. The box plots in Figure 3 shows the normalized saved running times for different n_{train} . As expected, the larger n_{train} , the more time is saved; setting $n_{\text{train}} = 15$ seems a fair compromise for the forthcoming experiments.

4.7 Base Realtime Approach

This is the main experiment where we evaluate the realtime approach. We run StreamLLM on all seven problems, comparing the static and adaptive approach and the two considered LLMs. We conducted the experiment twice to fathom the influence of randomness on the results. Figure 4 reports the average time reduction of both runs, where the time reduction is computed as $(\text{old runtime} - \text{new runtime}) / \text{old runtime}$.

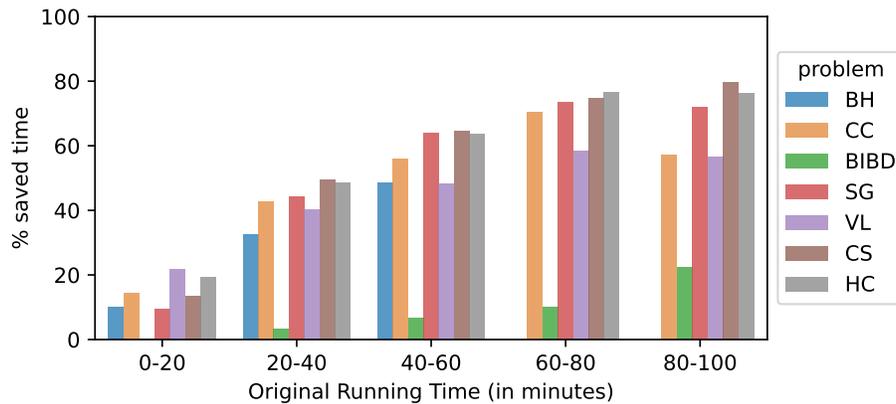


Fig. 5. Solving time reduction with respect to the original solving time when including streamliner generation time in the total streamlined running time.

Overall, the realtime approach achieves very encouraging results. Some runs achieve a reduction in running times of 98% and more. This is the case for the CS, the SG, and the HC problem. In some cases, the streamlined models could even solve most of the test instances in less than a second. The good performance on the HC problem suggests that the LLMs do not just copy and paste known streamliners from literature, but are also capable of dealing with new problems. For a few of the problems, most noticeably the BIBD problem, the reduction is only moderate. This is not so surprising considering that the model for BIBD already includes symmetry-breaking constraints and so it turns out to be more challenging to find significant time savings. It is worth noting that BIBD also sees very minor improvement in the approach by Spracklen et al. (Spracklen, Dang, et al. 2023).

Overall, Claude works slightly better than GPT. Comparing the adaptive and static approaches, there is no clear winner. This might be because the adaptive feedback reduces the exploratory potential of the LLMs.

Our analysis thus far compared the time reduction of streamlined models over the unstreamlined model, but did not consider the time StreamLLM spent on streamliner generation. To revisit the realtime scenario as sketched in the introduction, Figure 5 shows the percentage of saved time when including the streamliner generation time in the running time of the streamlined model³. As expected, the saved time for instances that take less than 20 minutes is relatively poor since more than half of the time is spent on streamliner generation. For larger instances, however, the generation time is relatively insignificant, and StreamLLM shows a remarkably strong performance on most of the problems, with the exception of the BIBD problem. Note that, the BH problem has no instances that take more than 60 minutes.

4.8 Obfuscation and Disguise

This experiment evaluates how obfuscation and disguise affect the performance of our StreamLLM approach. Figure 6 shows the results. It is interesting to see is that for GPT, our approach performs, on average, slightly better on disguised problems than on the original problems. For Claude, it performs slightly better on the original problems. In any case, the good results on the disguised problems suggest that the LLMs do not only copy and paste online available streamliners for known problems but can make sense of the underlying problem and can find streamliners for new problems. For the obfuscated problems, our approach performs, on average, worse than

³Since there are only a few instances where the unstreamlined model takes between 100 and 120 minutes, we disregarded this interval.

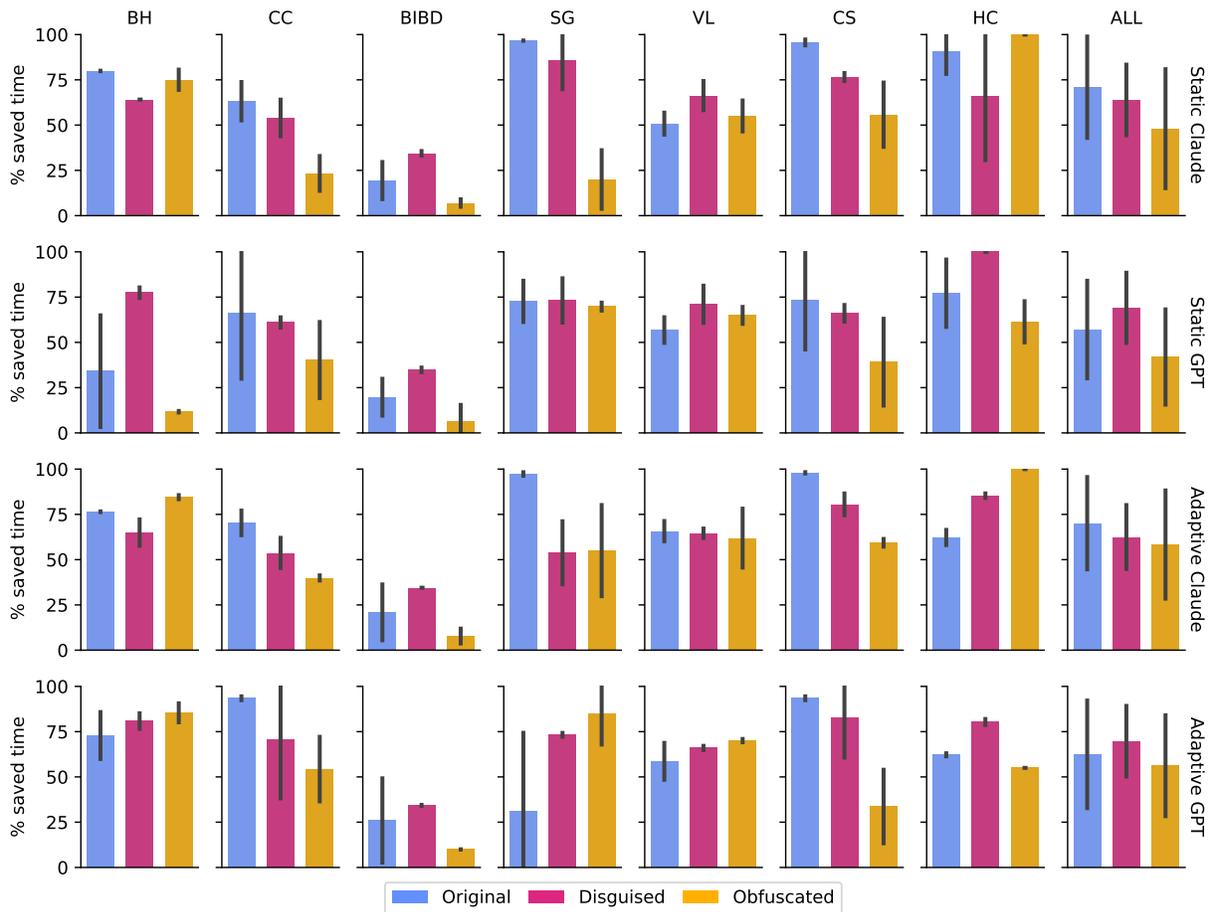


Fig. 6. Percentage of reduction in solving time for all seven original problems, as well as their disguised and obfuscated versions, with both static and adaptive approaches using LLMs Claude and GPT. The bars and black lines denote the means and standard deviations, respectively.

on the original and disguised problems. This suggests that LLMs prefer semantically richer language over pure abstract terminology to make better sense of variables and constraints.

4.9 Combinations

We evaluate streamliner combinations on the problems BH, CC, and VL, which we selected because they exhibited moderate performance improvements with individual streamliners, neither the dramatic gains seen in problems like SG nor the minimal improvements observed with BIBD. As shown in Figure 7, combining streamliners proves more effective than individual constraints for BH and VL while slightly reducing performance for CC.

4.10 Comparison with Previous Work

We replicate a subset of Spracklen et al.'s (Spracklen, Dang, et al. 2023) experiments on our benchmark instances to establish a baseline for comparing our base approach against. However, this comparison should be interpreted

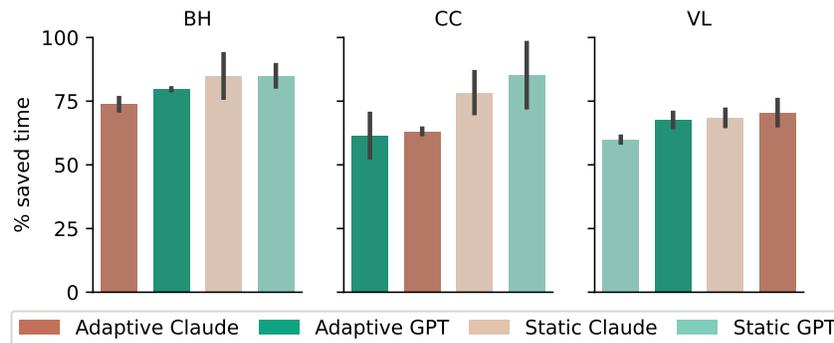


Fig. 7. Percentage of reduction in solving time obtained by combinations of streamliners with both static and adaptive approaches using LLMs Claude and GPT. The bars and black lines denote the means and standard deviations, respectively.

cautiously due to the following significant differences between the approaches: (i) the technical pipelines differ—their method uses Conjure and Essence for solving and modeling, whereas we employ MiniZinc for both purposes; (ii) the benchmark problems used in their experiments differ from our benchmark set.

We used the top-performing streamliners from the repository published by Spracklen et al.⁴ as the initial candidate pool and evaluated them on our benchmark instances. BIBD, CS, SG, and VL are the only problems that appear in both benchmarks. However, the Conjure model for VL is significantly different from our MiniZinc model, making it incompatible with our input instances. Thus, we only compare the remaining three problems. We first evaluated the initial candidates on our training instances to determine the three best streamliners for each problem, which we then evaluated on the test instances. The best triples saved 3%, 17%, and 21% time for the BIBD, CS, and SG problems, respectively. Only three of the nine streamliners from these triples could satisfy all the test instances.

We see two reasons for these savings being much lower than those reported by Spracklen et al.: (i) we consider all instances towards the total savings, while they only average over the improved instances; (ii) the difficulty of our benchmark instances was calibrated for MiniZinc, which need not necessarily align with the difficulty for Conjure. Finally, it is also worth noting, for pragmatic reasons, we compared their approach against our base approach, which operates in real time and does not generate and train streamliners offline like their method. Furthermore, their approach allows streamliners to be combinations of constraints, while our base approach considers only individual constraints.

4.11 Offline Approach

We conduct the offline approach on the BH and CC problems, which show a middle-ground performance in the main experiment, running each problem twice to assess the impact of randomness. Figure 8 presents the performance of the current best triple during training and its corresponding performance on test instances. The strong correlation between training and test instance trends validates our approach of inferring streamliner quality from smaller training instances. The performance graph shows significant quality improvements in the early stages before gradually leveling off. By the end of training, all runs achieve triples of streamliners that save over 90% of running time on test instances, demonstrating that this offline scenario with extended training times yields better results than the realtime approach. Among the top triples from each run, combinations make up slightly more than half of the streamliners, with Claude generating two-thirds and GPT-4 contributing one-third.

⁴<https://github.com/stacs-cp/automated-streamliner-portfolios>

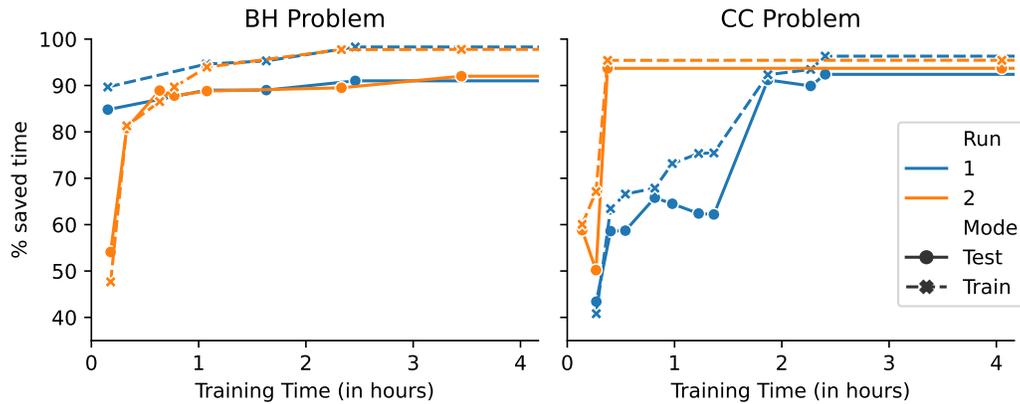


Fig. 8. Percentage of reduction in solving time by the offline approach on both training and test instances with respect to the length of the training phase for two runs each on two problems.

This distribution suggests that the diversity achieved through multiple prompts and LLMs enhances overall results.

5 Analysis of the Streamliners

In this section, we analyze the relationship between the individual streamliners within the best-performing triple for each problem. These triples are selected based on their performance in the main experiment outlined in Section 4.7. Table 1 shows the contribution of each of the streamliners and the original model, i.e., the percentage of test instances it performs best. The table also shows the percentage of unsatisfiable test instances for each streamliner. Further, one can see whether a streamliner is an implied constraint, a symmetry-breaking constraint, or does not preserve satisfiability.

It is interesting to see that for the different problems, the contribution of streamliners is distributed very differently. For the HC problem, one of the streamliners contributes to 98% of the instances, while the other two have almost no impact. In contrast to this, for the VL problem, the contribution of the three streamliners, as well as the original model, are distributed almost evenly. Further, one can see that for some of the problems, the original model has no contribution at all, which means that for each instance there is at least one of the streamliners in the triple that performs better than the original model.

It is worth noting that for some of the triples, the streamlined models make none or only a few of the instances unsatisfiable, whilst others have rather high percentages of unsatisfiable instances. Especially for the BH and the BIBD problem, even the streamliner with the highest contribution makes about one-third of the instances unsatisfiable. Of the 21 total streamliners mentioned in Table 1, 10 streamliners managed to satisfy all the test instances. We further manually inspected these 10 streamliners to find that 7 of them are indeed solution preserving, i.e., either implied or symmetry-breaking. All the 21 streamliners are listed in the appendix.

To complement this numerical analysis, we now present and discuss a few of the streamliners covered in the table.

Black Hole (BH).

```
s1: constraint x[26] == 26;
```

This streamliner for the Black Hole problem fixes the card with ID 26 to appear in the 26th position of the play sequence. It reduces the search space and allows about 60% of our instances to be solved more efficiently.

Table 1. Percentages of contribution and percentage of unsatisfiable instances for all streamliners constituting the best-performing triples for all problems. The type column indicates whether the streamliner is an implied constraint (i), a symmetry-breaking constraint (s), or a constraint that does not preserve satisfiability (-). A (u) denotes cases where the classification remains uncertain.

Problem	Contribution (%)				Unsat Instances (%)			Type		
	original	s1	s2	s3	s1	s2	s3	s1	s2	s3
BH	11	60	26	2	32	13	77	-	-	-
CC	0	85	12	4	2	2	44	-	-	-
BIBD	7	56	31	6	36	0	79	-	u	-
SG	0	74	25	2	0	5	0	s	-	s
VL	25	29	24	22	0	0	0	i	i	i
CS	0	84	12	3	0	5	0	i	-	u
HC	0	98	1	1	0	0	1	i	u	-

However, adding this constraint does not preserve satisfiability and makes about one-third of our instances unsatisfiable.

Social Golfers (SG).

s1: constraint forall(g in Golfer) (assign[g,2] = ((g-1) + n_per_group) mod n_groups + 1);
s3: constraint forall(g in Golfer) (assign[g,1] = (g-1) div n_per_group + 1);

The first shown symmetry-breaking constraint, s1, assigns each golfer to a specific group in week two, while s3 assigns each golfer to another specific group in week one. These constraints fix certain elements of the assign array and thereby narrow the search space. Note that s3 is a standard symmetry-breaking constraint that appears in online resources⁵. Because of the promising performance of s1, we had a closer look at it and tested whether it is also generalizable to larger instances. The results are remarkable: We came across many instances that could not be solved by the original model within 10 hours, but were solved by the streamlined model within 3 minutes. This corresponds to a runtime reduction of more than 99%.

Hypergraph Coloring (HC).

s1: constraint sum(c in Color) (num_vertices_of_color[c]) = num_vertices;

This constraint for the Hypergraph Coloring problem ensures that the sum of the vertices of every color equals the total number of vertices in the graph. Although it is already implied by other constraints in the encoding, its explicit inclusion significantly improves the solver performance.

6 Conclusion and Future Work

Our StreamLLM system and its analysis demonstrate that LLMs can effectively generate streamlining constraints for constraint satisfaction problems. Our experiments across seven diverse problems show impressive runtime reductions compared to unstreamlined models. The system can generate effective streamliners within minutes using only small training instances, making it practical even for problems with hour-long solving times. StreamLLM performs significantly better on the original and the disguised problem variants than on the obfuscated ones. This suggests a similarity to human experts, who naturally leverage the semantic context.

⁵https://www.csplib.org/Problems/prob010/models/social_golfers1.mzn.html

Our work gives rise to several promising research directions and possible extensions of our approach: (i) Knowledge distillation could compress the LLM components into smaller, locally deployable language models specifically tuned for streamliner generation. (ii) The observation that disguised problems sometimes yield better results suggests systematically exploring problem translations as a meta-optimization strategy. (iii) Developing LLM-supported formal verification methods for streamliner soundness could enable meaningful evaluation on unsatisfiable instances. This would allow us to infer that the original instance itself was unsatisfiable if the streamlined instance yields an unsatisfiable outcome. And finally, (iv) combining LLM-generated streamliners with evolutionary optimization techniques could improve their effectiveness while maintaining the rapid generation capabilities demonstrated in this work.

While this study focuses exclusively on constraint satisfaction problems, the general approach can also be applied to optimization problems. In our recent work (Voboril, Peruvemba Ramaswamy, et al. 2025), we applied the approach to optimization problems and achieved results surpassing the previously best-known solutions.

Acknowledgments

The authors thank Carlos Ansótegui for helpful discussions. This research is supported by the Austrian Science Funds (FWF), projects 10.55776/COE12, 10.55776/P36688, and 10.55776/P36420. Part of this work was carried out while the third author was a visiting researcher at the Simons Institute for the Theory of Computing, Berkeley.

References

- Ö. Akgün, A. M. Frisch, I. P. Gent, C. Jefferson, I. Miguel, and P. Nightingale. 2022. “Conjure: Automatic generation of constraint models from problem specifications.” *Artificial Intelligence*, 310, 103751.
- Ö. Akgün, A. M. Frisch, B. Hnich, C. Jefferson, and I. Miguel. 2011. “Conjure Revisited: Towards Automated Constraint Modelling.” *CoRR*, abs/1109.1774. <http://arxiv.org/abs/1109.1774> arXiv: 1109.1774.
- C. Ansótegui, F. Manyà, J. Ojeda, J. M. Salvia, and E. Torres. 2022. “Incomplete MaxSAT approaches for combinatorial testing.” *J. Heuristics*, 28, 4, 377–431. doi:10.1007/S10732-022-09495-3.
- R. L. Bras, C. P. Gomes, and B. Selman. 2014. “On the Erdős Discrepancy Problem.” In: *Principles and Practice of Constraint Programming - 20th International Conference, CP 2014, Lyon, France, September 8-12, 2014. Proceedings* (Lecture Notes in Computer Science). Ed. by B. O’Sullivan. Vol. 8656. Springer, Lyon, France, 440–448. doi:10.1007/978-3-319-10428-7_33.
- T. B. Brown et al.. 2020. “Language Models are Few-Shot Learners.” In: *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*. Ed. by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin. Vol. 33. Curran Associates, Inc., New York, NY, USA, 1877–1901. <https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfcb4967418bfb8ac142f64a-Abstract.html>.
- J. Charnley, S. Colton, and I. Miguel. 2006. “Automatic Generation of Implied Constraints.” In: *ECAI 2006, 17th European Conference on Artificial Intelligence*. IOS Press, Riva del Garda, Italy, 73–77.
- M. Chen et al.. 2021. “Evaluating Large Language Models Trained on Code.” *CoRR*, abs/2107.03374. <https://arxiv.org/abs/2107.03374> arXiv: 2107.03374.
- G. Chu and P. J. Stuckey. 2015. “Dominance breaking constraints.” *Constraints*, 20, 2, 155–182. doi:10.1007/s10601-014-9173-7.
- S. Colton and I. Miguel. 2001. “Constraint Generation via Automated Theory Formation.” In: *Principles and Practice of Constraint Programming - CP 2001, 7th International Conference*. Springer, Paphos, Cyprus, 575–579. doi:10.1007/3-540-45578-7_39.
- N. Dang, Ö. Akgün, J. Espasa, I. Miguel, and P. Nightingale. 2022. “A Framework for Generating Informative Benchmark Instances.” In: *28th International Conference on Principles and Practice of Constraint Programming, CP 2022, July 31 to August 8, 2022, Haifa, Israel (LIPICs)*. Ed. by C. Solnon. Vol. 235. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Haifa, Israel, 18:1–18:18. doi:10.4230/LIPICs.CP.2022.18.
- M. Díaz, R. L. Bras, and C. P. Gomes. 2017. “In Search of Balance: The Challenge of Generating Balanced Latin Rectangles.” In: *Integration of AI and OR Techniques in Constraint Programming - 14th International Conference, CPAIOR 2017, Padua, Italy, June 5-8, 2017, Proceedings* (Lecture Notes in Computer Science). Ed. by D. Salvagnin and M. Lombardi. Vol. 10335. Springer, Padua, Italy, 68–76. doi:10.1007/978-3-319-59776-8_6.
- J. K. Fichte, M. Hecher, and S. Szeider. 2020. “Breaking Symmetries with RootClique and LexTopSort.” In: *Principles and Practice of Constraint Programming - 26th International Conference, CP 2020, Louvain-la-Neuve, Belgium, September 7-11, 2020, Proceedings* (Lecture Notes in Computer Science). Ed. by H. Simonis. Vol. 12333. Springer, Louvain-la-Neuve, Belgium, 286–303. doi:10.1007/978-3-030-58475-7_17.
- P. Flener, A. M. Frisch, B. Hnich, Z. Kiziltan, I. Miguel, J. Pearson, and T. Walsh. 2001. “Symmetry in Matrix Models.” In: *Proceedings of SymCon*. Vol. 1. Citeseer, Paphos, Cyprus, 43–51.

- A. M. Frisch, B. Hnich, Z. Kiziltan, I. Miguel, and T. Walsh. 2009. "Filtering algorithms for the multiset ordering constraint." *Artificial Intelligence*, 173, 2, 299–328. doi:10.1016/j.artint.2008.10.007.
- A. M. Frisch, B. Hnich, Z. Kiziltan, I. Miguel, and T. Walsh. 2006. "Propagation algorithms for lexicographic ordering constraints." *Artificial Intelligence*, 170, 10, 803–834. doi:10.1016/j.artint.2006.05.001.
- A. M. Frisch, C. Jefferson, B. Martinez-Hernandez, and I. Miguel. 2007. "Symmetry in the Generation of Constraint Models." In: *Proceedings of the International Symmetry Conference*.
- A. M. Frisch, C. Jefferson, and I. Miguel. 2004. "Symmetry Breaking as a Prelude to Implied Constraints: A Constraint Modelling Pattern." In: *ECAI 2004, 16th European Conference on Artificial Intelligence*. IOS Press, Valencia, Spain, 171–175. doi:10.1007/978-3-540-30201-8_18.
- A. M. Frisch, I. Miguel, and T. Walsh. 2003. "CGRASS: A System for Transforming Constraint Satisfaction Problems." In: *Recent Advances in Constraints, Joint ERCIM/CoLogNET International Workshop on Constraint Solving and Constraint Logic Programming*. Springer, Cork, Ireland, 15–30. doi:10.1007/978-3-540-24662-6_2.
- A. M. Frisch, I. Miguel, and T. Walsh. 2001. "Symmetry and Implied Constraints in the Steel Mill Slab Design Problem." In: *Proceedings of the CP'01 Workshop on Modelling and Problem Formulation*. Springer, Paphos, Cyprus, 8–15.
- M. R. Garey and D. S. Johnson. 1979. *Computers and intractability*. Vol. 174. freeman San Francisco, San Francisco, CA, USA.
- C. P. Gomes, A. Sabharwal, and B. Selman. 2006. "Model Counting: A New Strategy for Obtaining Good Bounds." In: *Proceedings, The Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference, July 16-20, 2006, Boston, Massachusetts, USA*. AAAI Press, Boston, MA, USA, 54–61. <http://www.aaai.org/Library/AAAI/2006/aaai06-009.php>.
- C. P. Gomes and M. Sellmann. 2004. "Streamlined Constraint Reasoning." In: *Principles and Practice of Constraint Programming - CP 2004, 10th International Conference, CP 2004, Toronto, Canada, September 27 - October 1, 2004, Proceedings* (Lecture Notes in Computer Science). Ed. by M. Wallace. Vol. 3258. Springer Verlag, Toronto, Canada, 274–289. doi:10.1007/978-3-540-30201-8_22.
- A. Grover, T. Achim, and S. Ermon. 2018. "Streamlining Variational Inference for Constraint Satisfaction Problems." In: *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*. Ed. by S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett. Curran Associates, Montréal, Canada, 10579–10589. <https://proceedings.neurips.cc/paper/2018/hash/02ed812220b0705fabb868ddb17ea20-Abstract.html>.
- M. J. H. Heule, M. Kauers, and M. Seidl. 2019. "Local Search for Fast Matrix Multiplication." In: *Theory and Applications of Satisfiability Testing - SAT 2019 - 22nd International Conference, SAT 2019, Lisbon, Portugal, July 9-12, 2019, Proceedings* (Lecture Notes in Computer Science). Ed. by M. Janota and I. Lynce. Vol. 11628. Springer, Lisbon, Portugal, 155–163. doi:10.1007/978-3-030-24258-9_10.
- A. Itzhakov and M. Codish. 2022. "Complete symmetry breaking constraints for the class of uniquely Hamiltonian graphs." *Constraints*, 27, 1-2, 8–28. doi:10.1007/S10601-021-09323-8.
- G. Lample, T. Lacroix, M. Lachaux, A. Rodriguez, A. Hayat, T. Lavril, G. Ebner, and X. Martinet. 2022. "HyperTree Proof Search for Neural Theorem Proving." In: *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*. Ed. by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh. Vol. 35. Curran Associates, Inc., New Orleans, LA, USA, 26337–26349. http://papers.nips.cc/paper_files/paper/2022/hash/a8901c5e85fb8e1823bbf0f755053672-Abstract-Conference.html.
- R. LeBras, C. P. Gomes, and B. Selman. 2012. "From Streamlined Combinatorial Search to Efficient Constructive Procedures." In: *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, July 22-26, 2012, Toronto, Ontario, Canada*. Ed. by J. Hoffmann and B. Selman. AAAI Press, Toronto, Canada, 499–506. doi:10.1609/AAAI.V26I1.8147.
- J. H. M. Lee and A. Z. Zhong. 2022. "Exploiting Functional Constraints in Automatic Dominance Breaking for Constraint Optimization." In: *28th International Conference on Principles and Practice of Constraint Programming, CP 2022, July 31 to August 8, 2022, Haifa, Israel (LIPICs)*. Ed. by C. Solnon. Vol. 235. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Haifa, Israel, 31:1–31:17. doi:10.4230/LIPICs.CP.2022.31.
- J. Lee, T. Le, J. Chen, and D. Lee. 2023. "Do Language Models Plagiarize?" In: *Proceedings of the ACM Web Conference 2023, WWW 2023, Austin, TX, USA, 30 April 2023 - 4 May 2023*. Ed. by Y. Ding, J. Tang, J. F. Sequeda, L. Aroyo, C. Castillo, and G. Houben. ACM, Austin, TX, USA, 3637–3647. doi:10.1145/3543507.3583199.
- Z. Liu, L. Chew, and M. J. H. Heule. 2021. "Avoiding Monochromatic Rectangles Using Shift Patterns." In: *Proceedings of the Fourteenth International Symposium on Combinatorial Search, SOCS 2021, Virtual Conference [Jinan, China], July 26-30, 2021*. Ed. by H. Ma and I. Serina. AAAI Press, Jinan, China, 225–227. doi:10.1609/SOCS.V12I1.18591.
- L. Lovász. 1973. "Coverings and colorings of hypergraphs." In: *Proc. 4th Southeastern Conference of Combinatorics, Graph Theory, and Computing*. Utilitas Mathematica Publishing, 3–12.
- R. T. McCoy, S. Yao, D. Friedman, M. Hardy, and T. L. Griffiths. 2023. "Embers of Autoregression: Understanding Large Language Models Through the Problem They are Trained to Solve." *CoRR*, abs/2309.13638. arXiv: 2309.13638. doi:10.48550/ARXIV.2309.13638.
- S. Minaee, T. Mikolov, N. Nikzad, M. Chenaghlu, R. Socher, X. Amatriain, and J. Gao. 2024. "Large Language Models: A Survey." *CoRR*, abs/2402.06196. arXiv: 2402.06196. doi:10.48550/ARXIV.2402.06196.
- N. Nethercote, P. J. Stuckey, R. Becket, S. Brand, G. J. Duck, and G. Tack. 2007. "MiniZinc: Towards a standard CP modelling language." In: *International Conference on Principles and Practice of Constraint Programming* (Lecture Notes in Computer Science). Vol. 4741. Springer Verlag, Providence, RI, USA, 529–543. doi:10.1007/978-3-540-74970-7_38.

- K. Pei, D. Bieber, K. Shi, C. Sutton, and P. Yin. 2023. “Can Large Language Models Reason about Program Invariants?” In: *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA* (Proceedings of Machine Learning Research). Ed. by A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett. Vol. 202. PMLR, Honolulu, HI, USA, 27496–27520. <https://proceedings.mlr.press/v202/pei23a.html>.
- S. Polu, J. M. Han, K. Zheng, M. Baksys, I. Babuschkin, and I. Sutskever. 2023. “Formal Mathematics Statement Curriculum Learning.” In: *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, Kigali, Rwanda. <https://openreview.net/forum?id=-P7G-8dmSh4>.
- S. Prestwich and J. C. Beck. 2004. “Exploiting Dominance in Three Symmetric Problems.” In: *Fourth International Workshop on Symmetry and Constraint Satisfaction Problems*. Springer, Toronto, Canada, 63–70.
- B. Romera-Paredes et al. 2024. “Mathematical discoveries from program search with large language models.” *Nature*, 625, 7995, 468–475. doi:10.1038/S41586-023-06924-6.
- F. Rossi, P. Van Beek, and T. Walsh. 2006. *Handbook of constraint programming*. Elsevier, Amsterdam, Netherlands.
- A. Schutt, P. J. Stuckey, and A. R. Verden. 2011. “Optimal Carpet Cutting.” In: *Principles and Practice of Constraint Programming - CP 2011 - 17th International Conference, CP 2011, Perugia, Italy, September 12-16, 2011. Proceedings* (Lecture Notes in Computer Science). Ed. by J. H. Lee. Vol. 6876. Springer, Perugia, Italy, 69–84. doi:10.1007/978-3-642-23786-7_8.
- C. Smith, C. P. Gomes, and C. Fernández. 2005. “Streamlining Local Search for Spatially Balanced Latin Squares.” In: *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK, July 30 - August 5, 2005*. Ed. by L. P. Kaelbling and A. Saffiotti. Professional Book Center, Edinburgh, Scotland, 1539–1540. <http://ijcai.org/Proceedings/05/Papers/post-0460.pdf>.
- P. Spracklen, Ö. Akgün, and I. Miguel. 2018. “Automatic Generation and Selection of Streamlined Constraint Models via Monte Carlo Search on a Model Lattice.” In: *Principles and Practice of Constraint Programming - 24th International Conference, CP 2018, Lille, France, August 27-31, 2018, Proceedings* (Lecture Notes in Computer Science). Ed. by J. N. Hooker. Vol. 11008. Springer, Lille, France, 362–372. doi:10.1007/978-3-319-98334-9_24.
- P. Spracklen, N. Dang, Ö. Akgün, and I. Miguel. 2023. “Automated streamliner portfolios for constraint satisfaction problems.” *Artificial Intelligence*, 319, 103915. doi:10.1016/J.ARTINT.2023.103915.
- P. Spracklen, N. Dang, Ö. Akgün, and I. Miguel. 2019. “Automatic Streamlining for Constrained Optimisation.” In: *Principles and Practice of Constraint Programming - 25th International Conference, CP 2019, Stamford, CT, USA, September 30 - October 4, 2019, Proceedings* (Lecture Notes in Computer Science). Ed. by T. Schiex and S. de Givry. Vol. 11802. Springer, Stamford, CT, USA, 366–383. doi:10.1007/978-3-030-30048-7_22.
- P. Spracklen, N. Dang, Ö. Akgün, and I. Miguel. 2020. “Towards Portfolios of Streamlined Constraint Models: A Case Study with the Balanced Academic Curriculum Problem.” *CoRR*, abs/2009.10152. <https://arxiv.org/abs/2009.10152> arXiv: 2009.10152.
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. 2017. “Attention is All you Need.” In: *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*. Ed. by I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett. Curran Associates, Long Beach, CA, USA, 5998–6008. <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>.
- F. Voboril, V. P. R., and S. Szeider. Jan. 2025. *Code and Instances for the Paper: Generating Streamlining Constraints with Large Language Models*. (Jan. 2025). doi:10.5281/zenodo.14757597.
- F. Voboril, V. Peruvemba Ramaswamy, and S. Szeider. 2025. “Balancing Latin Rectangles with LLM-Generated Streamliners.” In: *31st International Conference on Principles and Practice of Constraint Programming (CP 2025)* (Leibniz International Proceedings in Informatics (LIPIcs)). Ed. by M. G. de la Banda. Vol. 340. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 36:1–36:17. ISBN: 978-3-95977-380-5. doi:10.4230/LIPIcs.CP.2025.36.
- F. Voboril, V. P. Ramaswamy, and S. Szeider. May 2025. “StreamLLM: Enhancing Constraint Programming with Large Language Model-Generated Streamliners.” English. In: *2025 IEEE/ACM 1st International Workshop on Neuro-Symbolic Software Engineering (NSE)*. IEEE Computer Society, (May 2025), 17–22. ISBN: 979-8-3315-1460-0. doi:10.1109/NSE66660.2025.00010.
- J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Richter, F. Xia, E. Chi, Q. V. Le, and D. Zhou. 2022. “Chain-of-Thought Prompting Elicits Reasoning in Large Language Models.” In: *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*. Ed. by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh. Vol. 35. Curran Associates, Inc., New Orleans, LA, USA, 24824–24837. https://proceedings.neurips.cc/paper_files/paper/2022/file/9d5609613524ecf4f15af0f7b31abca4-Paper-Conference.pdf.
- J. Wetter, Ö. Akgün, and I. Miguel. 2015. “Automatically Generating Streamlined Constraint Models with Essence and Conjure.” In: *Principles and Practice of Constraint Programming - 21st International Conference, CP 2015, Cork, Ireland, August 31 - September 4, 2015, Proceedings* (Lecture Notes in Computer Science). Ed. by G. Pesant. Vol. 9255. Springer, Cork, Ireland, 480–496. doi:10.1007/978-3-319-23219-5_34.
- H. Wu, C. Barrett, and N. Narodytska. 2023. “Lemur: Integrating Large Language Models in Automated Program Verification.” In: *The 3rd Workshop on Mathematical Reasoning and AI at NeurIPS’23*. New Orleans, LA, USA. <https://openreview.net/forum?id=NxHl2SPHYT>.

F. F. Xu, U. Alon, G. Neubig, and V. J. Hellendoorn. 2022. “A systematic evaluation of large language models of code.” In: *MAPS@PLDI 2022: 6th ACM SIGPLAN International Symposium on Machine Programming, San Diego, CA, USA, 13 June 2022*. Ed. by S. Chaudhuri and C. Sutton. ACM, San Diego, CA, USA, 1–10. doi:10.1145/3520312.3534862.

A Streamliners

Black Hole (BH).

```
s1: constraint x[26] == 26;
s2: constraint forall(i in 1..51) (x[i] != x[i+1] + 1);
s3: constraint x[52] == 52;
```

Carpet Cutting (CC).

```
s1: constraint forall(i in Stairs, j in st_rec_ids[i]) (st_rec_x[j] mod st_rec_len[j] = 0);
s2: constraint forall(i in Stairs, j in st_rec_ids[i]) (st_rec_y[j] mod st_rec_wid[j] = 0);
s3: constraint forall(i in Rooms, j in rm_rec_ids[i]) (rm_rec_x[j] mod MinRmRecSize = 0);
```

Balanced Incomplete Block Design (BIBD).

```
s1: constraint forall(i in rows)
    (sum(j in cols where j mod 2 = 1) (bool2int(m[i, j])) >= r div 2);
s2: constraint forall(i in rows) (sum(j in cols) (bool2int(m[i, j])) >= k div 2);
s3: constraint forall(i in rows, j in cols where i = j) (m[i, j] = false);
```

Social Golfers (SG).

```
s1: constraint forall(g in Golfer) (assign[g,2] = ((g-1) + n_per_group) mod n_groups + 1);
s2: constraint forall(w in Week, gr in Group) (assign[gr, w] = gr);
s3: constraint forall(g in Golfer) (assign[g,1] = (g-1) div n_per_group + 1);
```

Vessel Loading (VL).

```
s1: constraint forall(c in Containers) (Left[c] <= deck_width - width[c]);
s2: constraint max([Right[c] | c in Containers]) <= deck_width;
s3: constraint forall(c1, c2 in Containers where c1 < c2)
    (Left[c1] <= deck_width - width[c1] \\/ Left[c2] <= deck_width - width[c2]);
```

Car Sequencing (CS).

```
s1: constraint forall(i in options)
    (count(j in 1..n_cars) (b_seq_confs[i, j] = 1)
     = sum(c in configurations) (confs[c, i] * n_cars_by_confs[c]));
s2: constraint forall(i in 1..n_cars-1)
    (seq_confs[i] >= seq_confs[i+1]
     -> forall(j in options) (b_seq_confs[j, i] >= b_seq_confs[j, i+1]));
s3: constraint forall(i in 1..n_cars-2)
    (seq_confs[i] = seq_confs[i+1] -> seq_confs[i+1] != seq_confs[i+2]);
```

Hypergraph Coloring (HC).

```
s1: constraint sum(c in Color) (num_vertices_of_color[c]) = num_vertices;
s2: constraint forall(c in Color)
    (num_vertices_of_color[c] <= (num_vertices div num_colors) + (max_imbalance div 2));
s3: constraint forall(v1, v2 in Vertex where v1 < v2)
    (if exists(e in Hyperedge) (incidence[e, v1] = 1 /\ incidence[e, v2] = 1) then
```

```
coloring[v1] <= coloring[v2]  
endif);
```

Received 25 April 2025; accepted 24 October 2025