



Guided Tensor Lifting

YIXUAN LI, University of Edinburgh, United Kingdom

JOSÉ WESLEY DE SOUZA MAGALHÃES, University of Edinburgh, United Kingdom

ALEXANDER BRAUCKMANN, University of Edinburgh, United Kingdom

MICHAEL F. P. O'BOYLE, University of Edinburgh, United Kingdom

ELIZABETH POLGREEN, University of Edinburgh, United Kingdom

Domain-specific languages (DSLs) for machine learning are revolutionizing the speed and efficiency of machine learning workloads as they enable users easy access to high-performance compiler optimizations and accelerators. However, to take advantage of these capabilities, a user must first translate their legacy code from the language it is currently written in, into the new DSL. The process of automatically lifting code into these DSLs has been identified by several recent works, which propose program synthesis as a solution. However, synthesis is expensive and struggles to scale without carefully designed and hard-wired heuristics. In this paper, we present an approach for lifting that combines an enumerative synthesis approach with a Large Language Model used to *automatically* learn the domain-specific heuristics for program lifting, in the form of a probabilistic grammar. Our approach outperforms the state-of-the-art tools in this area, despite only using *learned* heuristics.

CCS Concepts: • **Computing methodologies** → **Probabilistic reasoning; Heuristic function construction**; • **Theory of computation** → **Program verification**; • **Software and its engineering** → **Search-based software engineering; Retargetable compilers; Source code generation**; Software maintenance tools.

Additional Key Words and Phrases: Program Synthesis, Lifting, Tensor Algebra, Code Optimization, Code Generation, Large Language Model

ACM Reference Format:

Yixuan Li, José Wesley de Souza Magalhães, Alexander Brauckmann, Michael F. P. O'Boyle, and Elizabeth Polgreen. 2025. Guided Tensor Lifting. *Proc. ACM Program. Lang.* 9, PLDI, Article 227 (June 2025), 23 pages. <https://doi.org/10.1145/3729330>

1 Introduction

Recent years have witnessed rapid growth in the number and importance of machine learning workloads. While used in a diverse number of applications, their fundamental building block is tensor contractions, which dominate execution time. For this reason, a large number of specialized tensor domain-specific languages (DSLs) have appeared, capable of producing high-performance code ([1, 6, 13, 14, 17, 34]). Their associated compilers are capable of extracting domain-specific information to exploit hardware-specific features and vendor-tuned libraries.

To access such performance, applications have to be written in one or more high-level DSLs. While this is acceptable for new applications, it means that existing workloads written in standard programming languages are unable to directly access a platform's potential performance. While

Authors' Contact Information: Yixuan Li, University of Edinburgh, Edinburgh, United Kingdom, yixuan.li.cs@ed.ac.uk; José Wesley de Souza Magalhães, University of Edinburgh, Edinburgh, United Kingdom, jwesley.magalhaes@ed.ac.uk; Alexander Brauckmann, University of Edinburgh, Edinburgh, United Kingdom, alexander.brauckmann@ed.ac.uk; Michael F. P. O'Boyle, University of Edinburgh, Edinburgh, United Kingdom, mob@inf.ed.ac.uk; Elizabeth Polgreen, University of Edinburgh, Edinburgh, United Kingdom, elizabeth.polgreen@ed.ac.uk.



This work is licensed under a Creative Commons Attribution 4.0 International License.

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

ACM 2475-1421/2025/6-ART227

<https://doi.org/10.1145/3729330>




Fig. 1. Overview of STAGG. We query the LLM to provide 10 possible solutions in TACO that are equivalent to the input code C. Based on the LLM response, we build a probabilistic grammar and enumerate the space of template programs described by the said grammar. We validate a candidate using I/O examples, and if it passes all tests, we proceed to verification to prove equivalence with the original C implementation. The input code is also analyzed to predict the dimensionality of the left-hand side tensor in the solution.

manually rewriting a program to a DSL may be a worthwhile cost, it becomes a serious impediment if it has to be repeated for new emerging DSLs. This problem of manual porting or *lifting* existing code to higher-level DSLs has been identified by several recent works that propose automated techniques. The most popular approaches use varying forms of program synthesis, where a DSLs space is searched for a matching program ([15, 25, 37]). However, program synthesis is expensive and struggles to scale to multi-dimensional tensor workloads.

To overcome this scalability issue, existing schemes rely on aggressive hard-wired heuristics that trade-off coverage for time. In C2TACO [25], domain-specific polyhedral analysis is used to prune the search space. This works well on low-dimensional problems but suffers from exponential growth. Similarly, in Tenspiller [35], the user provides a template to aid search. While narrowly effective, such heuristics limit portability and are a limit to generalization.

A completely different approach is to use neural machine translation based on large language models (LLMs). They have proved highly successful with a number of program generation tasks [8, 10]. They are fast and scale with program complexity, but unfortunately are inaccurate. What we would like is to combine the power of LLMs with the accuracy of synthesis.

This paper explores a novel combination of LLMs and program synthesis. It uses an LLM to suggest a number of possible solutions. It then builds a probabilistic grammar of templates, based on the proposed solutions, and then uses this grammar to drive an enumerative search of grammar templates. Our approach, termed STAGG (Synthesis of Tensor Algebra Guided by Grammars) is able to outperform all existing approaches. It achieves 99% lifting accuracy on a pre-existing large-scale benchmark suite of dense tensor algebra and is able to do this without any pre-wired heuristics.

Contributions: This paper makes the following contributions:

- Two novel synthesis algorithms that combine LLM guesses and program synthesis to scalably lift dense tensor code.
- A large-scale evaluation of state-of-the-art tensor program lifting.
- Greater coverage than existing techniques.

2 Motivation

The tensor DSL we target in our paper is TACO [17]. Whilst TACO may be best known for sparse computation, it also gives superior performance over dense code on multicores and GPUs. Although recent work has tackled *matching* certain sparse computation to specific high-performance APIs [12, 19], this paper focuses solely on *lifting* legacy dense tensor computation to a high-level programming language [16].

TACO syntax is based on Einstein summation (einsum) notation, a language for representing linear algebra operations using indexing expressions. The TACO compiler takes a tensor expression as input and generates highly optimized kernels. TACO-generated code exploits the parallel nature of both dense tensor algebra operations and multi-core/GPU architectures. It uses domain-specific knowledge to optimize and auto-parallelize, which related work has reported results in average speedups of $1.8\times$ and $24.1\times$ over the original program on CPUs and GPUs, respectively [25].

Einsum notation: TACO supports einsum notation, as do other frameworks including PyTorch [34], Halide [36]. While these alternatives may support a larger set of operations, targeting TACO allows a direct apples-to-apples comparison against prior work [25, 35].

Einsum expressions consist of a sequence of indexing variables, each one representing an iterator over a different tensor dimension. The traditional einsum notation expresses tensor multiplication and implicit summation on the indices that are absent in the output tensor. TACO uses an extended version of the original notation that also supports subtraction and division. Unlike other einsum-based frameworks such as the NumPy [14] einsum API, the tensors in TACO programs must be explicitly declared. Figure 5 shows the TACO grammar addressed in this paper.

Problem statement: Formally, given a legacy program p_s , written in a low-level language such as C, STAGG aims to find an equivalent program p_t written in TACO, that meets the specification

$$\forall \vec{x}. p_s(\vec{x}) = p_t(\vec{x}),$$

where \vec{x} is a vector of input arguments. That is, p_t produces the same output as p_s on all possible inputs.

Example: The synthesis task that STAGG solves is to synthesize a TACO program p_t such that $\forall \vec{x}. p_s(\vec{x}) = p_t(\vec{x})$, where \vec{x} is a vector of inputs, in this case `Mat1`, `Mat2`, `Result`. We now illustrate this approach on the example input C program shown in Figure 2.

Given this program, STAGG first queries a large language model to ask for a set of candidate solutions. The prompt template we use is shown below in Prompt 1. This gives us the set of candidate solutions shown in Response 1.

STAGG then learns a probabilistic context-free grammar that captures this set of solutions as templates.

We describe how we learn this grammar in Section 4.3. The grammar in Figure 3 shows probabilities for each production rule in parentheses. Each tensor and constant in the grammar is treated as a symbolic variable that can later be replaced when the template is instantiated.

We use a weighted A* search to explore the space of the grammar, inspired by work in the literature [21, 23], enhanced with penalty functions that penalize (partial or complete) templates that fail to adhere to syntactic constraints. When a complete template is found, this is passed to a template validator, which searches for all possible instantiations of the template and evaluates them against a set of input-output examples. A valid template, in this instance, would be the template `a(i) = b(i, j) * c(j)`. This is instantiated to the concrete program `Result(i) = Mat1(i, j) * Mat2(j)`. We compile this TACO program using the TACO compiler into C code, and check with bounded model checking that the two pieces of C code are equivalent.

```

1 void function(int N, int* Mat1, int* Mat2, int* Result){
2     int* p_m1;
3     int* p_m2;
4     int* p_t;
5     int i, f;
6     p_m1 = Mat1;
7     p_t = Result;
8     for (f = 0; f < N; f++) {
9         *p_t = 0;
10        p_m2 = &Mat2[0];
11        for (i = 0; i < N; i++)
12            *p_t += *p_m1++ * *p_m2++;
13        p_t++;
14    }
15 }

```

Fig. 2. A C implementation of $\sum_{i=0}^{N-1} \text{Mat1}(f \times N + i) \cdot \text{Mat2}(i)$. The result is a dot product between the f -th row of Mat1 and vector Mat2. The equivalent synthesized TACO expression is $a(i) = b(i, j) * c(j)$.

You are a scientific assistant that knows a lot about transpilation. Translate the following C code to an expression in the TACO tensor index notation. The expression must be valid as input to the taco compiler. Return a list with 10 possible expressions. Return the list and only the list, no explanations.

{the input C program}

Prompt 1. The prompt requesting 10 TACO expressions for a given C program. The temperature we use is 1.0, and the role is “You are a scientific assistant that knows a lot about transpilation”.

```

1 t(f) = m1(i, f) * m2(f)
2 Result(i) = Mat1(i,f)*Mat2(f)
3 Target(i) := Mat1(f,i) * Mat2(i)
4 Result(f) = sum(f, mat1(f, i) * mat2(i))

```

Response 1. LLM-generated candidate solutions for matrix product computations based on the implementation in Figure 2. Displayed are a subset of the 10 generated solutions, trimmed for brevity.

```

1 PROGRAM      ::= TENSOR1 "=" EXPR (1)
2 TENSOR1      ::= "a(i)" (1)
3 EXPR         ::= TENSOR (0) | CONSTANT (0) | EXPR OP EXPR (1)
4 OP           ::= "+" (0.2) | "-" (0) | "*" (0.8) | "/" (0)
5 TENSOR       ::= "b(i, j)" (0.2) | "b(j, k)" (0.1) | ...
               ::= "c(i)" (0.3) | "c(j)" (0.2) | "c(k)"

```

Fig. 3. A probabilistic context-free grammar template.

3 Overview of STAGG

Lifting to tensor DSLs is a challenging problem for program synthesis, and existing enumerative techniques are capable of accurate translation but rely on hand-written heuristics in order to scale. In contrast, highly scalable machine-learning-based approaches like language models fail to give accurate translations due to the complexity of the benchmarks. The key insight behind STAGG is that we can achieve the best of both worlds by using an LLM to *learn* the heuristics for an enumerative solver.

To that end, STAGG, as shown in Figure 1, implements a multi-staged hybrid synthesis approach.

- ① First, we construct a prompt based on the input C code and ask the LLM to propose 10 translated solutions.
- ② We proceed to construct a probabilistic grammar, which represents the space of solutions in the form of templates.
- ③ We then search this space of templates with a two-stage enumerative search: first, we search the space of templates with a search inspired by A^* , then, given a template, we search for a valid completion of the template against a set of input-output examples.
- ④ If a completed template is found that satisfies all input-output examples, we perform bounded verification with a bounded model checker to validate that the completed template is equivalent to the original C code. If it fails verification, we return to the template enumeration stage.

4 Learning a Grammar of Templates

The first step of STAGG uses a large language model to generate a set of candidate solutions for the synthesis problem in hand, using the prompt shown in Prompt 1. This gives us a set of candidate solutions, P . We ask for 10 solutions, but we parse in as many solutions as the LLM gives us (which is sometimes more than 10) and discard any syntactically incorrect solutions.

Given a set of incorrect candidate solutions from the LLM, we hypothesize that, even though none of the candidate solutions were precisely correct, the correct solution is likely to lie in the neighborhood of the LLM's guesses. To that end, we characterize this neighborhood using a probabilistic grammar of templates. We use a context-free grammar but note that, in principle, any probabilistic model that characterizes the neighborhood of guesses could be used. First, let us define some of the preliminaries we will need for this section.

4.1 Preliminaries

Definition 4.1 (Context-Free Grammar, CFG). A context-free grammar (CFG) is a 4-tuple,

$$G = \langle V, \Sigma, R, S \rangle.$$

V is a finite set of non-terminal symbols. Σ with $\Sigma \cap V = \emptyset$ is a finite set of terminal symbols. $R \subseteq V \times (V \cup \Sigma)^*$ is a finite set of production rules, where $*$ denotes the Kleene star. Each production rule in a context-free grammar is of the form $\alpha \rightarrow \beta$, where α is a symbol in the set of non-terminals V , and β is a string composed of symbols from $(V \cup \Sigma)^*$. $S \in V$ is the start symbol of the grammar G .

Given a context-free grammar $G = \langle V, \Sigma, R, S \rangle$, with $x, y \in (V \cup \Sigma)^*$ and a rule $(\alpha \rightarrow \beta) \in R$, we write $x \Rightarrow_G^* y$ if there exist strings $u, v \in (V \cup \Sigma)^*$ such that $x = u\alpha v$ and $y = u\beta v$, and we write $x \Rightarrow^* y$ if either $x = y$ or $x \Rightarrow x_1 \Rightarrow \dots \Rightarrow x_n \Rightarrow y$ for $n \geq 0$. The language generated by G , denoted \mathcal{L} , is the set of all strings over Σ that can be derived from the start symbol S by applying a sequence of production rules from R . Formally, we define

$$\mathcal{L} = \{\sigma \in \Sigma^* \mid S \Rightarrow_G^* \sigma\}.$$

Definition 4.2 (Weighted Context-Free Grammar, CFG). A weighted context-free grammar (wCFG) is a 5-tuple,

$$\text{wCFG} = \langle V, \Sigma, R, S, W \rangle,$$

contains a $G = \langle V, \Sigma, R, S \rangle$, and a weighted function $W : R \rightarrow \mathbb{R}^+$ maps each production rule to a positive real number.

Definition 4.3 (Probabilistic Context-Free Grammar, pCFG). A probabilistic context-free grammar,

$$\text{pCFG} = \langle V, \Sigma, R, S, \mathbb{P} \rangle,$$

is a specific instance of wCFG, with the additional constraint that the weights of the production rules expanding each non-terminal must sum to one. The probability function $\mathbb{P} : R \rightarrow [0, 1]$ in pCFG assigning a probability $\mathbb{P}(\alpha \rightarrow \beta)$ to each rule. The probabilities must satisfy

$$\forall \alpha \in V. \sum_{\beta: (\alpha \rightarrow \beta) \in R} \mathbb{P}(\alpha \rightarrow \beta) = 1.$$

4.2 Constructing a Grammar of Templates

Capturing the search space as a grammar of templates rather than complete TACO programs has two advantages: first, in comparison to using the full TACO grammar, it reduces the search space that the enumerative synthesis process has to search; and second, it allows us to group semantically equivalent but syntactically different candidate expressions together as one when constructing the weighted context-free grammar. For example, expressions like `1 t(f) = m1(i, f) * m2(f)` and `3 Target(i) := Mat1(f, i) * Mat2(i)` in LLM Response 1, are equivalent in structure (we use preprocessing to swap `:=` to `=` before parsing) shown in Figure 4, yet they would yield different terminal rules in the full grammar due to variations in notation.




Fig. 4. Expression standardization. We omit part of the derivation for brevity.

Given a set of candidate solutions, the first step is to construct a grammar of *templates* that captures the full set of solutions. The full grammar for TACO programs, G_{TACO} is shown in Figure 5. A TACO program is any program in $\mathcal{L}(G_{TACO})$.

Definition 4.4 (Templates). We define a TACO *template* τ to be any string obtained by taking a program $p \in \mathcal{L}(G_{TACO})$ and replacing all tensors with symbolic tensor variables, denoted t_1, t_2, \dots and all constants with constant symbols $const_1, const_2, \dots$.

Given a set of candidate programs from the LLM, P , we aim to find a grammar G_τ that contains a set of templates \mathcal{T} that allow us to generate all programs $p \in P$. A substitution $S = (t_1 \mapsto s_1^1, t_2 \mapsto s_2^1, \dots, const_1 \mapsto s_1^c)$ is a mapping from symbolic tensor variables and constant symbols to terminal symbols in the grammar G_{TACO} . A template *generates* a program if $\exists S. \tau.\{S\} = p$, where

$\tau.S$ indicates the result of replacing all occurrences of t_1 with s_1^t , and t_2 with s_2^t etc in t . Thus, our requirement on our grammar G is the following constraints: first that

$$\forall \tau. (\tau \in \mathcal{T}) \implies \tau \in \mathcal{L}(G),$$

and second that

$$\forall p \in P. \exists \tau \in \mathcal{T} \wedge \exists S. \tau.\{S\} = p.$$

This is obviously trivially satisfied by the complete grammar of TACO programs, yet it is undesirable as it has not reduced our search space. We also aim to create a grammar that is as small as possible, whilst avoiding over-fitting, and we attempt to optimize this trade-off by the method of construction described in the following section.

```

1 PROGRAM      ::= TENSOR "=" EXPR
2 TENSOR       ::= IDENTIFIER | IDENTIFIER "(" INDEX-EXPR ")"
3 EXPR         ::= TENSOR | CONSTANT | "(" EXPR ")" | "-" EXPR |
                  EXPR "+" EXPR | EXPR "-" EXPR |
                  EXPR "*" EXPR | EXPR "/" EXPR
4 INDEX-EXPR   ::= INDEX-VAR | INDEX-VAR "," INDEX-EXPR
5 INDEX-VAR    ::= "i" | "j" | "k" | "l"
6 IDENTIFIER   ::= LETTER ( LETTER | INTEGER ) *
7 CONSTANT     ::= INTEGER
8 INTEGER      ::= DIGIT +
9 LETTER       ::= "a" | "b" | ... | "z" | "A" | "B" | ... | "Z"
10 DIGIT       ::= "0" | "1" | "2" | ... | "9"

```

Fig. 5. The grammar for TACO expression in Extended Backus–Naur form, defining the syntax for tensor expressions, identifiers, constants, and basic arithmetic expressions. The $*$ symbol denotes Kleene star, indicates zero or more repetitions of the preceding element, while $+$ denotes Kleene plus, requires one or more occurrences.

4.2.1 Templatized Candidate Solution. We obtain the grammar G_τ by first inferring a template for each solution in P . The first step involves parsing each $p \in P$ into an Abstract Syntax Tree (AST), a structured representation that captures the hierarchical organization of the expression. The AST organizes operations, tensor identifiers, and indices as distinct nodes, allowing systematic traversal and manipulation. For example, consider the expression `1 t(f) = m1(i, f) * m2(f)` in Response 1 can be parsed as the left tree in Figure 4. We then transform the AST in three stages: *Tensor Templatization*, *Index Standardization*, and *Constant Templatization*.

Tensor Templatization. We replace each tensor name in the expression with a symbolic tensor variable. From hereon, we use a, b, c , as the symbolic tensor variables t_1, t_2, t_3, \dots , to align with the variable names in the code examples. The identifiers are assigned in alphabetical order—starting with a for the left-hand side tensor and using b, c, d, \dots sequentially for tensors on the right-hand side, based on their order of first appearance. The expression `t(f) = m1(i, f) * m2(f)` will be transformed into `a(f) = b(i, f) * c(f)` by this step.

Index Standardization. The index standardization step ensures that each tensor expression in the grammar uses a consistent set of index variables, irrespective of the original indices in the input expression. Each unique index variable encountered in an expression is mapped to the next available symbol from the canonical set $\{i, j, k, l\}$ in alphabetical order. The expression `a(f) =`

$b(i, f) * c(f)$ will be transformed to $a(i) = b(j, i) * c(i)$ by this step, as shown in the middle in Figure 4. The index variables do not need to be replaced by template instantiation as they are local variables to the program.

Constant Templatzation. Any constants in the candidate solutions are replaced with a symbolic constant `Const`. The template instantiation step will instantiate from a list of constants found in the input source code.

4.2.2 Refining the Grammar. Given the templatzed candidate solutions \mathcal{T} , we wish to construct a probabilistic grammar that represents the space of these solutions, without substantially over-fitting. The first step is to construct a context-free grammar that defines this set of templates.

We start with the base grammar of TACO programs, shown in Figure 5. We then restrict the set of tensor names to be the names we have chosen as symbolic tensor variables and constants, namely a, b, c, \dots and `const`, and also limit the set of index variables to be i, j, k, l, \dots . In theory, this permits 26 tensor IDs and 4 index variables, because one can always infer whether a variable is an index or a tensor identifier by context. In practice, we never need this many, and searching a space that includes up to 26 4-dimensional tensors is obviously impractical. This section addresses how we initially reduced this search space. Namely, by predicting the dimensions of the tensors in order to refine the grammar.

4.2.3 Predicting Tensor Dimensions. To accurately predict tensor dimensions for a given program, we combine insights from a language model (LLM) with static code analysis. Static analysis is used to predict the left-hand side (LHS) tensor dimensions of an expression, while the LLM is used to predict dimensions for the right-hand side (RHS). Static analysis, by analyzing the source code, can determine precisely the dimensions for the LHS tensor, but cannot do the same for the RHS, so we fall back on heuristics learned by the LLM for the RHS.

Definition 4.5 (Dimension list). We define a dimension list L_τ to be a list of integers (d_1, d_2, d_3, \dots) where d_i is the dimension for the i^{th} unique tensor in the template τ . We use $L[i]$ to indicate accessing the i^{th} element of the dimension list L , and $|L|$ to indicate the length of the list L . We list the dimensions of constants and variables as 0.

For example, the list $[1, 2, 4]$ indicates that the first tensor in the expression has 1 dimension, the second tensor has 2 dimensions, and the third tensor has 4 dimensions.

Dimension Prediction for RHS Tensors Using LLM. Given a set of templatzed solutions $\{\tau_1, \dots, \tau_{10}\} \in \mathcal{T}$, generated by the LLM, we compute the dimension list for each candidate solution, giving a set of lists $L_{\mathcal{T}} = \{L_{\tau_1}, \dots, L_{\tau_{10}}\}$. We then filter this list by length, and remove any list that has a length less than the maximum length, giving the filtered set

$$L_{\mathcal{T}}^* = \{l \in L_{\mathcal{T}} \mid |l| \geq |L_{\tau_i}| \forall L_{\tau_i} \in L_{\mathcal{T}}\}.$$

Finally, we return the list that appears most frequently in the filtered set, i.e.,

$$L_\tau = \arg \max_l |l \in L_{\mathcal{T}}^*|.$$

From here on, we denote this predicted list L , and refer to this as the predicted dimension list.

Integrating Static Analysis for LHS Tensors. We use static program analysis to examine the original program AST and predict the LHS dimension. We apply a dataflow analysis to recover the dimensions in the array accesses to recover the original dimensionality. For standard array accesses, e.g., $a(i, j)$, we simply count the number of variables used to index the base pointer. However, it is common that C programs access multi-dimensional elements using affine linear

expressions on index variables. In such cases, we use array delinearization [30] to recover the standard array access form and predict the dimensionality by counting the number of indexing variables. Additionally, some applications use explicit pointer arithmetic to iterate over arrays. We implement array recovery [11] to retrieve array access expressions from pointers and then apply delinearization and analyze the indexing expression. In case the output variable is not accessed through any memory indexing operation, we assume it is a scalar and predict zero-dimensionality.

As the left-hand side tensor necessarily appears first in the expression, we replace $L[1]$ with the predicted dimension for the first tensor from the static analysis.

4.2.4 Generating the Context-Free Grammar. Given a dimension list, we wish to generate a grammar that ensures we only enumerate combinations of indices required to make all possible tensor expressions that match the predicted dimensions (or, at least reduce this space as far as possible without increasing the complexity of the grammar significantly). For instance, if the dimension list is $[0, 1, 3]$, and at least one of the predicted solutions is $a = b(i) + c(i, j, k)$, we will modify the grammar to fix the production rule for the first tensor to restricted to a , and ensure the grammar can enumerate $b(i)$ and $c(i, j, k)$, $c(i, k, j)$, $c(j, i, k)$, $c(k, i, j)$, $c(j, k, i)$, $c(k, j, i)$ for any remaining tensor that appear in the expression. This permits all possible combinations of indexing of the tensors of the predicted dimensions, allowing the index used for $b(i)$ to be repeated in any position when indexing c .

Formally, we define this grammar generator as a set of constraints reasoning over the predicted dimension list L and the set of templates \mathcal{T} , which, if true, indicate that the production rule should be included in the grammar. We use $[c]r_i$ to denote that a production r_i is included within the grammar if c is true. $i(P)$ denotes the number of unique index variables in the set of programs \mathcal{T} . Rules without a constraint automatically appear in the grammar.

```

1      PROGRAM ::= TENSOR1 "=" EXPR
2      EXPR    ::= TENSOR | EXPR OP EXPR
3      OP      ::= "+" | "-" | "*" | "/"
4  [L[1] = 0] TENSOR1 ::= "a"
5  [L[1] = 1] TENSOR1 ::= "a(i)"
6  [L[1] = 2] TENSOR1 ::= "a(i, j)"
7      ...
8  [L[2] = 0] TENSOR ::= "b" | "Const"
9  [L[2] = 1] TENSOR ::= "b(i)" | [i(P) > 1] "b(j)" | [i(P) > 2] "b(k)" |
10     ...
11 [L[2] = 2] TENSOR ::= "b(i, j)" | "b(j, i)" | [i(P) > 2] "b(i, k)" | ...
12     ...
13 [L[3] = 0] TENSOR ::= "c" | "Const"
14 [L[3] = 1] TENSOR ::= "c(i)" | [i(P) > 1] "c(j)" | [i(P) > 2] "c(k)" | ...
15     ...

```

For every element in the dimension list, we add a new tensor id, a , b , c . . . , and index it with the number of index variables that correspond to the element in the dimension list. For an element i , where $L[i] = n$, and for a set of candidate programs where $i(\mathcal{T}) = m$, we add a production rule for every possible way of choosing n indices from m index variables. We then remove any production rules that could not be used parsing any candidate $\tau \in \mathcal{T}$, e.g., if no template in \mathcal{T} contains a 2-dimension tensor indexed with the same index variable twice, we will remove $b(i, i)$. An example generated template grammar is shown in Figure 6.

```

1 PROGRAM    ::= TENSOR1 "=" EXPR
2 TENSOR1    ::= "a(i)"
3 EXPR       ::= TENSOR | CONSTANT | EXPR OP EXPR
4 TENSOR     ::= "b(i, j)" | "b(j, i)" | "b(i, k)" | ...
               "c(i)" | "c(j)" | "c(k)"
5 CONSTANT   ::= "Const"
6 OP         ::= "+" | "-" | "*" | "/"

```

Fig. 6. An example generated template grammar, for the dimension list $[1, 2, 1, 0]$, with a maximum of 3 unique indices appearing in the candidate LLM solutions.

4.3 Assigning Probabilities to the Grammar

We now wish to assign a probability to each production rule in the grammar of TACO templates according to their frequency in the left-most derivations of the candidate solutions. Let us define a derivation as follows:

Definition 4.6 (Derivations). Given a context-free grammar G as previously defined, and a sentence $s \in \mathcal{L}(G)$, the derivation of s from S is a sequence of rules such that $S \xrightarrow{r_0} s_1 \xrightarrow{r_1} \dots s_n \xrightarrow{r_n} s$ and $r_0 \dots r_n \in R$. We denote the derivation of s by the sequence of rules r_0, \dots, r_n as $D_s = \{r_0, \dots, r_n\}$. The left-most derivation is a derivation such that all rules expand the left-most non-terminal symbol in the sentential form.

Given a set of templated solutions $\mathcal{T} \in \mathcal{L}(G_{\text{template}})$, we calculate a weight for each rule $r_i \in R$ as the number of times that rule appears in the left-most derivations of the programs. That is,

$$W[r_i] = \sum_{\tau_i \in \mathcal{T}} |r_i| \in D_{\tau_i},$$

where $|r_i|$ is the number of times r_i appears in the derivation. These weights reflect the usage frequency of each tensor with specific indices in the expressions provided by the LLM. Note that, for any production rules used to replace the tensor nonterminal symbols, e.g., **1DTENSOR**, which do not appear in any of the candidate solutions, we assign a default weight of 1. This assignment ensures that these combinations are considered during the synthesis process with a lower priority.

Using the weights calculated for both operators and tensors, we construct the corresponding probabilistic Context-Free Grammar ($pCFG_{\tau}$) by normalizing the weights into probabilities. For each non-terminal symbol α , the probability of applying the production rule $\alpha \rightarrow \beta$ is given by:

$$\mathbb{P}[\alpha \rightarrow \beta] = \frac{W[\alpha \rightarrow \beta]}{\sum_{\gamma} W[\alpha \rightarrow \gamma]},$$

where the sum in the denominator is over all production rules where α appears on the left-hand side, and γ may be any string of nonterminal and terminal symbols.

5 Searching the Template Space

We present two algorithms for searching the space of TACO templates. The first is based on a weighted A^* search in the literature [21, 23], which searches the grammar of TACO templates in a top-down manner. We extend this algorithm to incorporate a penalty score that accounts for known syntactic constraints on the target solution. The second is an adapted version of A^* , which combines bottom-up search with A^* heuristics.

5.1 Top-Down Weighted A*

Algorithm 1 outlines the top-down weighted A* search with penalties. The search operates over a probabilistic Context-Free Grammar derived from large language model (LLM) outputs. It maintains a queue of partial templates, represented as abstract syntax trees, which we can think of as the frontier of its search. This initially contains just the start symbol of the grammar. At each iteration, it must determine which of the partial templates should be further explored, which it does based on the cost of the path taken to reach those partial templates and an estimate of the cost required to extend the path all the way to the goal. The goal is, ultimately to find a complete template that we believe is likely to satisfy the specification, that is a template can generate a program p_t such that $\forall \vec{x}. p_t(\vec{x}) = p_s(\vec{x})$.

Thus, when choosing which partial template to further explore, the algorithm chooses the template with the minimum $f(x)$, where x is the partial template, defined as:

$$f(x) = c(x) + g(x) + X(x),$$

where $c(x)$ calculates the accumulated cost from the start node S to current node x . $g(x)$ is the heuristic estimate of the minimal cost to complete the expression from x to a goal node (a template we believe is likely to satisfy the specification) and $X(x)$ is a penalty term for expressions that violate domain-specific syntactic constraints. These are calculated as follows:

The accumulated cost $c(x)$ is calculated as the sum of the costs of the production rules applied along the path to x :

$$c(x) = \sum_{r_i \in D_x} -\log_2 \mathbb{P}[r_i],$$

where D_x is the sequence of production rules used to reach the node x and $\mathbb{P}[r_i]$ is the probability of production rule r_i . This cost function transforms probabilities into additive costs, suitable for the A* search.

The heuristic function $g(x)$ estimates the minimal additional cost required to complete the partial expression at node x to a full expression. It is defined as:

$$g(x) = \begin{cases} 0 & \text{if } x \in \Sigma^*, \\ -\sum_{x_i \in V} \log_2 h(x_i) & \text{otherwise,} \end{cases}$$

where x_i are the non-terminal symbols in the partial expression x , and $h(\alpha)$ is the maximal probability of deriving any terminal string from non-terminal α . The value $h(\alpha)$ is defined recursively for each non-terminal α :

$$\forall \alpha \in V, \quad h(\alpha) = \max_{\alpha \rightarrow \beta \in R} \left(\mathbb{P}[\alpha \rightarrow \beta] \times \prod_{\beta_i \in \beta} h(\beta_i) \right),$$

with the base case $h(\alpha) = 1$, if $\alpha \in \Sigma$. This equation represents the maximal probability of deriving a terminal string from α , accounting for the probabilities of production rules and the maximal probabilities of its components.

The penalty function $X(x)$ assigns additional costs to expressions that do not meet specific domain criteria. This function can be formalized as follows:

$$X(x) = \begin{cases} \sum_{a \in A} X_a(x) & \text{if } x \text{ violates criterion } a, \\ 0 & \text{otherwise,} \end{cases}$$

There are 5 criteria $\{a_1, \dots, a_5\} \in A$, and their penalty scores are defined as follows (note that an infinite penalty score effectively means these expressions will never be considered):

- $\mathcal{X}_{a_1}(x) = 10$, where a_1 is violated if the grammar includes a constant expression, the length of x exceeds 3, and x either 1) contains fewer than 2 tensors with index i or 2) lacks a constant expression. This penalty biases the search against expressions with multiple tensors but inadequate index variety or missing constants.
- $\mathcal{X}_{a_2}(x) = 100$, where a_2 is violated iff x is not the same length as the length of the dimension list.
- $\mathcal{X}_{a_3}(x) = \infty$, where a_3 is violated if the tensor symbols in x are not in alphabetical order by order of first appearance. This penalty rule avoids enumerating templates that are structurally identical, and therefore can be instantiated into identical sets of programs.
- $\mathcal{X}_{a_4}(x) = \infty$, where a_4 is violated if x is a complete template (no non-terminal symbols), and repeatedly applies addition, subtraction, or division operations on the same tensor.
- $\mathcal{X}_{a_5}(x) = \infty$, where a_5 is violated if x is a complete template (no non-terminal symbols), and employs fewer than half of the operations defined in the grammar.

Algorithm 1 Top-Down Enumerator

```

1: procedure ENUMERATE( $pCFG_\tau$ )
2:    $Q \leftarrow \{(0, pCFG_\tau.S)\}$                                  $\triangleright$  Initialize queue with start symbol of grammar
3:   while  $Q \neq \emptyset$  do
4:      $(f, x) \leftarrow Q.pop()$                                  $\triangleright$  Remove template with minimal  $f$ 
5:     if  $\text{depth}(x) > \text{maxDepth}$  then
6:       continue                                               $\triangleright$  Skip if maximum depth exceeded
7:     if  $x \in \Sigma^*$  then                                      $\triangleright$  If no non-terminals remain in  $x$ 
8:        $S \leftarrow \text{VALIDATE}(x)$                                 $\triangleright$  Try to instantiate  $x$ 
9:       if  $S \neq \perp$  then
10:        if  $\text{VERIFY}(x.\{S\})$  then
11:          return  $x.\{S\}$ 
12:        for  $x'. \text{ s.t. } (x \xrightarrow{r} x' \wedge r \in pCFG_\tau.R)$  do       $\triangleright$  Iterate over all possible expansions of  $x$ 
13:           $Q \leftarrow Q \cup \{c(x') + g(x') + \mathcal{X}(x'), x'\}$ 
14:   return Failure                                              $\triangleright$  Return Failure if no valid expression is found
  
```

Search: the search is shown in Algorithm 1. The algorithm keeps a queue of expressions in a queue, which you can consider to be the frontier of the search. At each iteration, it selects the expressions with the lowest total score f from the queue. If the expression is a complete expression, i.e., it contains only terminal symbols from the grammar, we then send this to the validation procedure described in Section 6. If the expression is a partial expression, the leftmost non-terminal of the expression is expanded according to all applicable production rules in the grammar, creating a new template for each production rule. These new expressions are all added to the queue and the process is repeated.

We set a depth limit of 6, and if any expression exceeds this depth, it is discarded. We calculate depth as the depth of the maximum child in the abstract syntax tree, excluding index expressions, e.g., $b(i)$ and $c(i, j)$ are both expressions of depth 1, and $b(i) + c(i, j)$ is an expression of depth 2.

5.2 Bottom-Up Weighted A*

The algorithm presented in the previous section takes a top-down approach to enumerating through the search space. This has advantages over bottom-up search, namely that it is known to find longer

programs faster than bottom-up search, which is biased towards shorter programs. Nevertheless, recent work has shown that guided bottom-up search can produce promising results [3]. To that end, we develop a new A* inspired bottom-up search algorithm, which we term bottom-up A*. The bottom-up search, shown in Algorithm 2 constructs expressions incrementally by starting with basic tensors and systematically combining them using operators, following a probabilistic context-free grammar. Again, the algorithm maintains a queue of expressions, and uses the same combination of cost, estimated cost to reach a goal state and the penalty function to determine which expression to expand first.

One key difference in the bottom-up search is the way we generate the template grammar. For the bottom-up search where the production rules only permit extending an expression by adding an operator and a new tensor to the end, effectively forcing the algorithm to enumerate programs shortest first. The grammar generator, given a predicted dimension list L , and a function $i(\mathcal{T})$ which calculates the number of unique indices in \mathcal{T} , is shown below:

```

1      PROGRAM    ::=  TENSOR1 "=" EXPR
2      EXPR       ::=  TENSOR2 TAIL1
3      TAIL1      ::=  ε | [|L| > 2] OP TENSOR3 TAIL2
4      TAIL2      ::=  ε | [|L| > 3] OP TENSOR4 TAIL3
5      ...
6      OP         ::=  "+" | "-" | "*" | "/"
7      [L[1] = 0] TENSOR1 ::= "a"
8      [L[1] = 1] TENSOR1 ::= "a(i)"
9      [L[1] = 2] TENSOR1 ::= "a(i, j)"
10     [L[2] = 0] TENSOR2 ::= "b" | "Const"
11     [L[2] = 1] TENSOR2 ::= "b(i)" | [i(P) > 1] "b(j)" | [i(P) > 2] "b(k)"
12     [L[2] = 2] TENSOR2 ::= "b(i, j)" | "b(i, j)" | "b(j, i)" |
13                           [i(P) > 2] "b(i, k)" | ...
14     [L[3] = 0] TENSOR3 ::= "c" | "Const"
15     [L[3] = 1] TENSOR3 ::= "c(i)" | [i(P) > 1] "c(j)" | [i(P) > 2] "c(k)" | ...
16     ...

```

An example of generated grammar is shown in Figure 7. Weights and probabilities over the grammar are then calculated as described in Section 4.3.

```

1  PROGRAM    ::=  TENSOR1 "=" EXPR
2  TENSOR1    ::=  "a"
3  EXPR       ::=  1DTENSOR TAIL1
4  TAIL1      ::=  ε | OP 2DTENSOR TAIL2
5  TAIL2      ::=  ε | OP 1DTENSOR
6  2DTENSOR   ::=  "c(i, j)" | "c(j, i)" | "c(i, k)" | ...
7  1DTENSOR   ::=  "b(i)" | "d(k)"

```

Fig. 7. An example generated template grammar, for the dimension list $[0, 1, 2, 1]$, with a maximum of 3 unique indices appearing in the candidate LLM solutions. The rules for each tensor index expression include all possible permutations of indices. The symbol ϵ denotes the empty string.

The search algorithm maintains a queue of partial programs, as with the top-down search, which is initialized with the start symbol from the grammar. At each iteration, the program with

Algorithm 2 Bottom-Up Enumerator

```

1: procedure ENUMERATE( $pCFG_\tau, L$ )
2:    $Q \leftarrow \{(0, pCFG_\tau.S)\}$  ▷ Priority queue initialized with start node
3:   while  $Q \neq \emptyset$  do
4:      $(f, x) \leftarrow Q.pop()$  ▷ Remove template with minimal  $f$ 
5:     if  $|tensors(x)| = |L|$  then ▷ If number of tensors in  $x$  is the predicted number
6:       if  $x \notin \Sigma^*$  then
7:          $x \leftarrow \text{RemoveTail}(x)$  ▷ Remove any tail nonterminal symbol
8:          $S \leftarrow \text{VALIDATE}(x)$  ▷ Try to instantiate  $x$ 
9:         if  $S \neq \perp$  then
10:          if  $\text{VERIFY}(x.\{S\})$  then
11:            return  $x.\{S\}$  ▷ Return instantiated template
12:          for  $x'. \text{s.t. } (x \xrightarrow{r} x' \wedge r \in pCFG_\tau.R)$  do ▷ Iterate over all possible expansions of  $x$ 
13:             $Q \leftarrow Q \cup \{c(x') + g(x') + X(x'), x'\}$ 
14:  return Failure ▷ Return Failure if no valid expression is found

```

the minimum cost function, as before, is popped from the queue, expanded, and all the resulting programs are added to the queue.

The total cost function $f(x)$ for each partial expression x , is again defined as: $f(x) = c(x) + g(x) + X(x)$, where $c(x)$ is calculated as before. In the bottom-up search, we use a simplified estimate of the cost to complete the program, $g(x)$, which is defined as:

$$g(x) = \sum_{i=k}^{|L|} m(L[i+1]),$$

where k is the current number of tensors in x , L is the predicted dimension list, and $m(d)$ is the minimal cost to add a tensor of dimension d . The minimal cost $m(d)$ is computed as follows, where $\text{Tensors}(d)$ is the list of tensors in the grammar of dimension d , and $\mathcal{P}[t]$ is the maximum probability of any production rule in the grammar which adds the tensor of dimension d :

$$m(d) = -\log_2 \left(\max_{t \in \text{Tensors}(d)} \mathcal{P}[t] \right).$$

The penalty function is calculated as before, but with the criteria $\{b_1, b_2\} \in B$ defined as:

- $\mathcal{X}_{b_1}(x) = 100$, where b_1 is violated if the tensor symbols in x are not in alphabetical order by order of first appearance.
- $\mathcal{X}_{b_2}(x) = \infty$, where b_2 is violated if x contains at least as many tensors as predicted by the dimension list, and it uses fewer than half the operations available in the grammar, and 0 otherwise.

The bottom-up search uses fewer penalty criteria than the top-down search because the construction of the grammar encapsulates a number of these criteria already (for instance, the tensors are enumerated by predicted dimension list order).

The main difference between the top-down and the bottom-up search is that the bottom-up grammar is generated in a way that, at each intermediate step, a complete program can be inferred from the partial program and checked against the specification. Every time an expression is dequeued from the queue, if the expression contains a tail nonterminal symbol, e.g., `TAIL1`, `TAIL1`; we can remove the tail nonterminal symbol to give a complete template (i.e., a template that contains no non-terminal symbols). We can then return this template to the template validator. If it fails

validation, we will re-append the nonterminal symbol to the end of the expression and generate new expressions by expanding the non-terminal using all applicable production rules. The new expressions are added into the queue.

6 Validation

Once the synthesizer produces a complete template τ , we wish to check whether it can generate a program p_t that satisfies the requirement $\forall \vec{x}. p_s(\vec{x}) = p_t(\vec{x})$. Since checking this universally quantified formula is expensive, we first generate a set of tests in the form of input-output examples. This set of examples $\langle I, O \rangle$ is a list of input-output pairs where $I = \langle x_1 \mapsto v_1, \dots, x_n \mapsto v_n \rangle$ is a map that binds the τ input arguments $\vec{x} = (x_1, \dots, x_n)$ to concrete values (v_1, \dots, v_n) randomly generated and $O = (o_1, \dots, o_n)$ is the corresponding output produced when we execute p_s on the elements from I .

The TACO templates generated during the synthesis phase contain symbolic placeholders for tensors and constants. To validate a candidate τ we build a set $S = (s_1, \dots, s_m)$ of substitutions $s \mapsto x$ that map the symbolic symbols s in τ to input arguments x . We iterate through all possible permutations of S , where tensor symbols are mapped to concrete tensor inputs, and constants are mapped to a set of constants $C = (c_1, \dots, c_m)$ containing the constant values present in the source code of p_s . We can then generate a concrete program $p_t = \tau.\{S\}$, and execute p_t on the input-output examples in $\langle I, O \rangle$. If any instantiated concrete program satisfies all the input-output examples, we return this to the next stage of verification. We use S to assign concrete values to TACO symbols and run P_T to check its output. When building S , we rule out invalid substitutions based on the type of arguments and TACO symbols. More specifically, we discard substitutions that try to bind tensor symbols with dimension > 1 to scalars and vice versa.

If the validator succeeds, it returns the valid substitution S , if not it returns \perp to indicate there was no valid substitution.

We explore all possible valid substitutions until we find a substitution s^* that satisfies ϕ . This validation process returns a tuple $\langle P_T, s^* \rangle$ that is given as input to the verification phase.

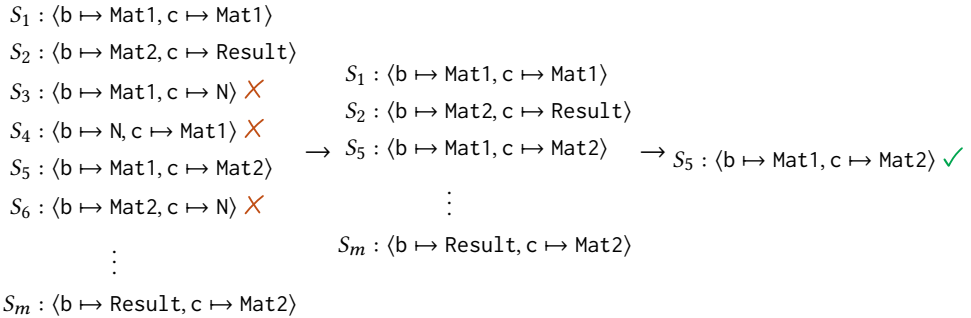


Fig. 8. A set of possible substitutions for the TACO program $a(i) = b(i, j) * c(j)$ and the inputs from the legacy program in Figure 2. We discard invalid substitutions and try the valid ones until we find one that satisfies the specification.

Example 6.1. Figure 8 shows a subset of the substitutions set S given the program in Figure 2 and the TACO candidate P_T produced by the synthesizer, $a(i) = b(i, j) * c(j)$. Each substitution binds a symbol in the right-hand side of P_T , i.e., b and c to one of the inputs of function. The substitutions with a \times mark next to it are invalid, since they contain unsound bindings. For

example, substitution S_3 binds c , a 1-dimensional tensor to N , which is a scalar. Such substitutions are discarded and the valid ones are tested to run the program until we find one in which P_T satisfies the specification. In this example, the correct substitution is S_5 , which binds b and c to arguments Mat1 and Mat2 respectively.

7 Verifier

We verify the correctness of a synthesized TACO program using bounded model checking. We compile both the original C and the TACO program to a common language within the MLIR compiler infrastructure [20]. Given a TACO program T and substitution S returned by the validator, we create NumPy code based on the indexing expressions of T and replace its variables for the concrete values specified in s^* . If the model checker fails to verify equivalence with the tuple $\langle T, s^* \rangle$, we return to the validation step and keep exploring different substitutions until we find one that satisfies the specification and passes verification. We then use the JAX compiler [6] to lower the NumPy code to MLIR.

From the MLIR files, we automatically generate C programs that create non-deterministic inputs, execute the original C and TACO code on copies of those inputs, and assert that the outputs are identical. We give this C program as input to CBMC [18], a bounded model checker for C, that verifies said assertion holds for all possible inputs up to a certain bound.

Floating-point equivalence is both challenging to verify and, in many cases, undesirable. For instance, many compiler optimizations simply do not preserve floating-point optimizations, in order to achieve runtime speed-ups. For this reason, we extend CBMC to support rational datatypes, and verify equivalence using rational datatypes.

8 Evaluation

To evaluate STAGG and its various components, we compare its performance against several established techniques on a diverse suite of queries. The query set includes 10 artificial examples and 67 real-world problems (61 derived from codebases reported in the literature [25] and 6 from the C++ based inference code of Llama [24].)

STAGG is implemented using an extended version of CBMC 6.3.1 with cvc5 version 1.0.5 as the underlying SMT solver. To generate initial candidate solutions, we use GPT-4 with the temperature set to 1.0. A timeout of 60 minutes is applied to each query. All experiments are conducted on a system equipped with an 11th Gen Intel® Core™ i5-1135G7 processor, 16 GiB of RAM, and running Ubuntu® 22.04.5 LTS. Additional configuration details, including grammar refinements and penalty modifications, are provided in the subsequent sections.

We compare the following approaches: (**STAGG^{TD}**) our approach, using the top-down A^* search described in Section 5.1; (**STAGG^{BU}**) our approach, using the bottom-up search described in Section 5.2; (**C2TACO**) An enumerative synthesis tool for lifting C to TACO code [25]. We compare to C2TACO both with and without the domain-specific heuristics; (**Tenspiller**) An enumerative synthesis tool based on the verified lifting framework [35]; (**LLM only**) A baseline approach that employs a large language model (GPT-4) to directly generate candidate solutions without additional heuristic-driven refinement or search. In addition, we perform ablation studies to evaluate the contribution of several components of STAGG. Namely, the grammar refinement; the probabilities of the grammar; and the penalty functions.

We aim to answer the following research questions:

RQ1: How does the performance of STAGG compare to the state-of-the-art enumerative synthesis tools?

RQ2: How does the choice of top-down or bottom-up A* search contribute to the performance of STAGG?

RQ3: How much do the penalty functions contribute to the performance of STAGG?

RQ4: How much does the grammar refinement contribute to the performance of STAGG?

RQ5: How much do the probabilities on the grammar contribute to the performance of STAGG?




Fig. 9. Cactus plot showing the number of benchmarks solved (x-axis) vs. time (y-axis, logarithmic) on the 67 real-world benchmarks. Each line corresponds to a different synthesizer, and the point at which each line indicates how many benchmarks the synthesizer solved before the time.

RQ1: Comparison of performance of STAGG to the state-of-the-art solvers: Figure 9 depicts the cumulative time each method takes over the 67 real-world benchmarks, and Figure 10 shows the success rate of different techniques. We compare STAGG to C2TACO on the full set of 77 benchmarks in Table 1. STAGG^{TD} solves 76 benchmarks, compared to C2TACO which solves 67. STAGG^{TD} solves all the benchmarks that C2TACO can solve, with an average solving time of 3.19s, compared to C2TACO’s 21.15s. STAGG^{BU} solves 73 benchmarks, and solves 66/67 of the benchmarks that C2TACO solves, with an average solving time of 2.11s on the mutually solved benchmarks. C2TACO without the domain heuristics enabled is significantly slower.

We are only able to run Tenspiler on the 67 real-world benchmarks, where it solves 52. STAGG^{TD} solves all 52 benchmarks that Tenspiler can solve, with an average time of 3.45s compared to Tenspiler’s average time of 4.56s. STAGG^{BU} solves only 50/52 of the benchmarks that Tenspiler can solve, but with an average time of 2.03s. This comprehensively answers RQ1: STAGG outperforms the state-of-the-art solvers, both in terms of coverage and speed.

RQ2: Performance Comparison of Top-Down vs Bottom-Up Search. Our results show that, while STAGG^{TD} solves more benchmarks than STAGG^{BU}, STAGG^{BU} is faster on commonly solved benchmarks (for queries solved by both, STAGG^{BU} achieves a lower average solving time (98.81 seconds) compared to STAGG^{TD} (108.83 seconds)), and it enumerates fewer candidates. It has one big disadvantage though, which is that it can only expand expressions by appending to the previous expression, rather than by expanding nodes on the left-hand side of the AST. In particular, this means it cannot solve benchmarks that require expressions with more balanced Abstract Syntax Trees or benchmarks that contain parentheses.




Fig. 10. Success rates of different approaches on the set of 67 real-world benchmarks.

Table 1. Comparison of benchmark-solving performance across different methods: The table reports the number of benchmarks solved (#), average solving time (time in seconds), and attempts across various benchmarks. The benchmarks are categorized into real-world benchmarks (67 in total), real-world + artificial benchmarks (77 in total), benchmarks solved by C2TACO, and benchmarks solved by Tenspiller. STAGG^{TD} and STAGG^{BU} demonstrate superior solving capabilities, solving more benchmarks overall compared to C2TACO and Tenspiller, with STAGG^{BU} achieving the fastest solving times for benchmarks solvable by C2TACO and Tenspiller. The results highlight the efficacy of STAGG over existing methods.

Methods	Real-World (67)		Real-World + Artificial (77)			Solved by C2TACO		Solved by Tenspiller	
	#	time	#	time	attempts	#	time	#	time
STAGG ^{TD}	66	121.88	76	106.13	44.55	67	3.19	52	3.45
STAGG ^{BU}	63	113.86	73	98.81	35.62	66	2.11	50	2.03
LLM	24	2.61	34	2.59	1.62	31	2.57	20	2.72
C2TACO	59	22.57	67	21.15	18.45	67	21.15	50	23.69
C2TACO.NoHeuristics	59	43.08	67	49.41	48.81	67	49.41	50	43.76
Tenspiller	52	4.56						52	4.56




Fig. 11. Impact of different grammar configurations in STAGG on success rates across all 77 benchmarks.

RQ3: Contribution of the penalty functions. Table 2 shows the decline in performance for both STAGG^{TD} and STAGG^{BU} approaches when individual penalty rules are removed. As each penalty rule is dropped, the number of queries solved decreases, highlighting the importance of these rules in achieving high query-solving efficiency.

Table 2. Impact of penalty rules on performance over 77 benchmarks (real-world + artificial). The table compares the number of benchmarks solved (#), the percentage of benchmarks solved (%), and the average solving time (time in seconds) for various configurations of STAGG. Removing penalty rules (e.g., Drop(A), Drop(B)) reduces the number of solved benchmarks and influences solving times. While STAGG^{TD} and STAGG^{BU} achieve high solving rates with the full penalty rules, dropping specific penalties often results in faster solving times but at the cost of reduced solving capability, as it failed solving complex benchmarks.

Methods	Real-World + Artificial (77)		
	#	%	time
STAGG ^{TD}	76	98.70%	106.13
STAGG ^{TD} .Drop(A)	71	92.21%	7.21
STAGG ^{TD} .Drop(a1)	72	93.51%	79.24
STAGG ^{TD} .Drop(a2)	75	97.40%	91.66
STAGG ^{TD} .Drop(a3)	72	93.51%	21.01
STAGG ^{TD} .Drop(a4)	75	97.40%	90.58
STAGG ^{TD} .Drop(a5)	75	97.40%	83.34
STAGG ^{BU}	73	94.81%	98.81
STAGG ^{BU} .Drop(B)	70	90.91%	68.18
STAGG ^{BU} .Drop(b1)	71	92.21%	48.95
STAGG ^{BU} .Drop(b2)	70	90.91%	68.75

Table 3. Performance comparison of different methods and grammar configurations over 77 benchmarks (real-world + artificial). The table shows the number of benchmarks solved (#), the percentage of benchmarks solved (%), the average solving time (time in seconds), and the number of synthesis attempts. STAGG^{TD} and STAGG^{BU} outperform C2TACO variants in solving more benchmarks. Variations of STAGG demonstrate the impact of grammar refinement, where dropping penalty rules (Drop(A), Drop(B)) or using alternative configurations (e.g., EqualProbability, LLMGrammar) affect the solving capability, time, and attempts.

Methods	Real-World + Artificial (77)			
	#	%	time	attempts
STAGG ^{TD}	76	98.70%	106.13	44.55
STAGG ^{TD} .Drop(A)	71	92.21%	7.21	13.65
STAGG ^{TD} .EqualProbability	73	94.81%	28.14	37.27
STAGG ^{TD} .LLMGrammar	52	67.53%	3.77	5.25
STAGG ^{TD} .FullGrammar	69	89.61%	91.15	874.29
STAGG ^{BU}	73	94.81%	98.81	35.62
STAGG ^{BU} .Drop(B)	70	90.91%	68.18	10.07
STAGG ^{BU} .EqualProbability	74	96.10%	180.31	62.78
STAGG ^{BU} .LLMGrammar	52	67.53%	2.74	2.60
STAGG ^{BU} .FullGrammar	68	88.31%	96.57	259.35
LLM	34	44.16%	2.59	1.62
C2TACO	67	87.01%	21.15	18.45
C2TACO.NoHeuristics	67	87.01%	49.41	48.81

RQ4 and RQ5: Contribution of grammar refinement and probabilities. Figure 11, 12 and Table 3 show the performance of difference configurations of STAGG: EqualProbability uses the refined grammar but replaces all probabilities in the generated pCFG with equal probabilities; FullGrammar uses the full TACO grammar in Figure 5 with equal probabilities; LLMGrammar uses the full TACO




Fig. 12. The performance of different configurations of STAGG on all 77 benchmarks.

grammar in Figure 5 with probabilities learned from the LLM responses. All these configurations use the penalty functions. Thus, in order to compare the contribution of the grammar refinement, we can compare the performance of LLMGrammar, which uses learned probabilities but no refinement, to STAGG. Dropping the refinement, here, results in us solving 31% fewer benchmarks. If we compare FullGrammar to EqualProbability, we can see that grammar refinement has less impact when learned probabilities are not used, but still results in a significant number of benchmarks being dropped.

In order to compare the contribution of the probabilities, we can compare the performance of EqualProbability to STAGG, where we note that using equal probabilities on the refined grammar results in an increase in the number of benchmarks solved for STAGG^{TD} , and an increase in solving speed for STAGG^{BU} , although neither is as impactful as the grammar refinement. In fact, the comparison between FullGrammar and LLMGrammar demonstrates that learned probabilities can have a negative impact if they are used in a grammar that is not general enough. Thus, our answer to RQ4 and RQ5 is that grammar refinement in combination with probabilities has a bigger impact on performance than either component part, but the refinement alone is more powerful than the probabilities alone.

9 Related Work

9.1 LLM-Guided Synthesis

One of the strategies to aid the search process in the program synthesis tool is using a neural network to guide the search. Seminal work by Balog et al. [2] synthesized array manipulation programs from I/O examples using a feedforward neural network (FNN) to build a probabilistic distribution over the target language. During the search, the synthesis algorithm expands partial programs based on the probabilities predicted by the FNN for the given I/O specification. Neural-guided synthesis has also been applied to solve string manipulation tasks by [31, 38], inductive logic programming [38], dataframe [4] and tensor [28, 37] processing, and code transpilation [26] using distinct models. SketchAdapt [29] uses a model to produce a program sketch as a starting point and completes said sketch through symbolic enumeration. Closer to our technique is Euphony [22],

a system that learns a probabilistic high-order grammar model from examples and searches for programs using A^* . All those techniques require the model to be trained on domain-specific data, which can be costly and an impediment for such techniques to be expanded to other tasks.

Pre-trained Large Language Models (LLMs) have been used to guide enumerative synthesis in class program synthesis tasks: Li et al. [23] developed two LLM-guided synthesis approaches to generate programs from logical formulas. The first method infers a probabilistic grammar using an LLM and enumerates said grammar using top-down A^* search; while the second queries the LLM to complete incorrect candidates produced by the enumeration phase. HYSYNTH [3] also derives a probabilistic grammar from LLM responses, but it uses bottom-up enumeration. LLMs are also used directly for verified lifting in LLMlift [5], where lifting of programs is performed leveraging GPT-4 [33] to guess candidate solutions and loop invariants to prove equivalence, with feedback being given to the LLM to correct any mistakes. This approach is remarkably effective but relies on the LLM being able to fix its own mistakes based on the feedback, which may not be trivial in complex domains [32].

9.2 Tensor Code Lifting

Driven by enormous advances in Machine Learning, there has been increasing interest in lifting tensor code to optimized targets. C2TACO [25] is a synthesis tool that generates TACO code from I/O examples. It implements a bottom-up enumerative algorithm, and it uses code analysis to restrict the search space of programs. mlirSynth [7] has a similar approach, but it lifts tensor programs across different MLIR dialects. In both methods, correctness is asserted using only I/O testing while STAGG performs bounded model checking to verify that the lifted programs are equivalent to their original counterpart. A different synthesis method was used in Tenspiller [35], which employs symbolic synthesis to generate programs in six different tensor DSLs. Tenspiller builds verification conditions and loop invariants to prove that the lifted program is equivalent to the original one. Unlike STAGG, which learns how to explore the search space in a fully automated way, all those techniques require hard-coded heuristics to make the search space tractable.

Another approach to lift code is API matching, in which source code is replaced by optimized library routines to improve performance. Examples in the tensor domain include KernelFaRer [9], which focuses on general matrix multiplication (GEMM), and ATC [27], which targets both GEMM and convolutions. SpEQ [19] introduces a method of translating sparse linear algebra codes to optimized targets using equality saturation applied to LLVM IR. However, these approaches are often tailored to specific APIs and are not portable. STAGG leverages the great learning capabilities of Large Language Models to infer the search space, which makes our technique extensible to different targets and to more unrestricted back-ends such as DSLs.

10 Conclusions

This paper presented STAGG, a novel approach that combines LLMs and program synthesis to lift legacy tensor code to DSLs. We use a set of LLM responses to infer a probabilistic context-free grammar that drives an enumerative search over the space of possible solutions. Our technique successfully lifts 99% of a large suite of benchmarks with an average lifting time of 3.19 seconds, outperforming existing state-of-the-art lifters in terms of coverage and synthesis time. Additionally, STAGG is able to automatically learn a search space, and it does not rely on any pre-defined heuristics. Future work will focus on expanding our technique to application domains other than tensor computation.

References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mane, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viegas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. [doi:10.48550/arXiv.1603.04467](https://arxiv.org/abs/1603.04467) [cs.DC]
- [2] Matej Balog, Alexander L. Gaunt, Marc Brockschmidt, Sebastian Nowozin, and Daniel Tarlow. 2017. DeepCoder: Learning to Write Programs. [arXiv:1611.01989](https://arxiv.org/abs/1611.01989) [cs.LG] <https://arxiv.org/abs/1611.01989>
- [3] Shraddha Barke, Emmanuel Anaya Gonzalez, Saketh Ram Kasibatla, Taylor Berg-Kirkpatrick, and Nadia Polikarpova. 2024. HYSYNTH: Context-Free LLM Approximation for Guiding Program Synthesis. *arXiv preprint arXiv:2405.15880* (2024). [doi:10.48550/arXiv.2405.15880](https://arxiv.org/abs/2405.15880)
- [4] Rohan Bavishi, Caroline Lemieux, Roy Fox, Koushik Sen, and Ion Stoica. 2019. AutoPandas: neural-backed generators for program synthesis. *Proc. ACM Program. Lang.* 3, OOPSLA, Article 168 (Oct. 2019), 27 pages. [doi:10.1145/3360594](https://doi.org/10.1145/3360594)
- [5] Sahil Bhatia, Jie Qiu, Sanjit A Seshia, and Alvin Cheung. 2024. Can LLMs Perform Verified Lifting of Code? *Technical Report* (2024). <https://www2.eecs.berkeley.edu/Pubs/TechRpts/2024/EECS-2024-11.pdf>
- [6] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Nécule, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. 2018. *JAX: composable transformations of Python+NumPy programs*. [http://github.com/google/jax](https://github.com/google/jax)
- [7] Alexander Brauckmann, Elizabeth Polgreen, Tobias Grosser, and Michael F. P. O'Boyle. 2024. mlirSynth: Automatic, Retargetable Program Raising in Multi-Level IR using Program Synthesis. In *Proceedings of the 32nd International Conference on Parallel Architectures and Compilation Techniques* (Vienna, AE, Austria) (PACT '23). IEEE Press, 39–50. [doi:10.1109/PACT58117.2023.00012](https://doi.org/10.1109/PACT58117.2023.00012)
- [8] Xinyun Chen, Chang Liu, and Dawn Song. 2018. Tree-to-Tree Neural Networks for Program Translation. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems* (Montréal, Canada) (NIPS'18). Curran Associates Inc., Red Hook, NY, USA, 2552–2562. <https://dl.acm.org/doi/10.5555/3327144.3327180>
- [9] João P. L. De Carvalho, Braedy Kuzma, Ivan Korostelev, José Nelson Amaral, Christopher Barton, José Moreira, and Guido Araújo. 2021. KernelFaRer: Replacing Native-Code Idioms with High-Performance Library Calls. *ACM Trans. Archit. Code Optim.* 18, 3, Article 38 (June 2021), 22 pages. [doi:10.1145/3459010](https://doi.org/10.1145/3459010)
- [10] Mehdi Drissi, Olivia Watkins, Aditya Khant, Vivaswat Ojha, Pedro Sandoval, Rakia Segev, Eric Weiner, and Robert Keller. 2018. Program Language Translation Using a Grammar-Driven Tree-to-Tree Model. [arXiv:1807.01784](https://arxiv.org/abs/1807.01784) [cs.LG] <https://arxiv.org/abs/1807.01784>
- [11] Björn Franke and Michael O'boyle. 2003. Array recovery and high-level transformations for DSP applications. *ACM Trans. Embed. Comput. Syst.* 2, 2 (May 2003), 132–162. [doi:10.1145/643470.643472](https://doi.org/10.1145/643470.643472)
- [12] Philip Ginsbach, Bruce Collie, and Michael F. P. O'Boyle. 2020. Automatically harnessing sparse acceleration. In *Proceedings of the 29th International Conference on Compiler Construction*. 179–190. [doi:10.1145/3377555.3377893](https://doi.org/10.1145/3377555.3377893)
- [13] Awni Hannun, Jagrit Digani, Angelos Katharopoulos, and Ronan Collobert. 2023. *MLX: Efficient and flexible machine learning on Apple silicon*. <https://github.com/ml-explore>
- [14] Charles R Harris, K Jarrod Millman, Stéfan J Van Der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. 2020. Array programming with NumPy. *Nature* 585, 7825 (2020), 357–362. [doi:10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2)
- [15] Shoaib Kamil, Alvin Cheung, Shachar Itzhaky, and Armando Solar-Lezama. 2016. Verified lifting of stencil computations. *SIGPLAN Not.* 51, 6 (2016), 711–726. [doi:10.1145/2980983.2908117](https://doi.org/10.1145/2980983.2908117)
- [16] Fredrik Kjolstad, Stephen Chou, David Lugato, Shoaib Kamil, and Saman Amarasinghe. 2017. Taco: A tool to generate tensor algebra kernels. In *2017 32nd IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 943–948. [doi:10.1109/ASE.2017.8115709](https://doi.org/10.1109/ASE.2017.8115709)
- [17] Fredrik Kjolstad, Shoaib Kamil, Stephen Chou, David Lugato, and Saman Amarasinghe. 2017. The tensor algebra compiler. *Proc. ACM Program. Lang.* 1, OOPSLA, Article 77 (Oct. 2017), 29 pages. [doi:10.1145/3133901](https://doi.org/10.1145/3133901)
- [18] Daniel Kroening and Michael Tautschnig. 2014. CBMC – C Bounded Model Checker. In *Tools and Algorithms for the Construction and Analysis of Systems*, Erika Ábrahám and Klaus Havelund (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 389–391. [doi:10.1007/978-3-642-54862-8_26](https://doi.org/10.1007/978-3-642-54862-8_26)
- [19] Avery Laird, Bangtian Liu, Nikolaj Børner, and Maryam Mehri Dehnavi. 2024. SpEQ: Translation of Sparse Codes using Equivalences. *Proc. ACM Program. Lang.* 8, PLDI, Article 215 (June 2024), 24 pages. [doi:10.1145/3656445](https://doi.org/10.1145/3656445)
- [20] Chris Lattner, Mehdi Amini, Uday Bondhugula, Albert Cohen, Andy Davis, Jacques Pienaar, River Riddle, Tatiana Shpeisman, Nicolas Vasilache, and Oleksandr Zinenko. 2020. MLIR: A Compiler Infrastructure for the End of Moore's Law. [arXiv:2002.11054](https://arxiv.org/abs/2002.11054) [cs.PL] <https://arxiv.org/abs/2002.11054>

- [21] Woosuk Lee, Kihong Heo, Rajeev Alur, and Mayur Naik. 2018. Accelerating search-based program synthesis using learned probabilistic models. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Philadelphia, PA, USA) (PLDI 2018). Association for Computing Machinery, New York, NY, USA, 436–449. doi:10.1145/3192366.3192410
- [22] Woosuk Lee, Kihong Heo, Rajeev Alur, and Mayur Naik. 2018. Accelerating search-based program synthesis using learned probabilistic models. *SIGPLAN Not.* 53, 4 (June 2018), 436–449. doi:10.1145/3296979.3192410
- [23] Yixuan Li, Julian Parsert, and Elizabeth Polgreen. 2024. Guiding Enumerative Program Synthesis with Large Language Models. In *Computer Aided Verification* (Montreal, Canada) (CAV 2024). Springer Nature Switzerland, Cham, 280–301. doi:10.1007/978-3-031-65630-9_15
- [24] llama 2024. llamacpp. <https://github.com/leloykun/llama2.cpp/>. Accessed: 2024-01-19.
- [25] José Wesley de Souza Magalhães, Jackson Woodruff, Elizabeth Polgreen, and Michael F. P. O’Boyle. 2023. C2TACO: Lifting Tensor Code to TACO. In *Proceedings of the 22nd ACM SIGPLAN International Conference on Generative Programming: Concepts and Experiences* (Cascais, Portugal) (GPCE 2023). Association for Computing Machinery, New York, NY, USA, 42–56. doi:10.1145/3624007.3624053
- [26] Benjamin Mariano, Yanju Chen, Yu Feng, Greg Durrett, and İsil Dillig. 2022. Automated transpilation of imperative to functional code using neural-guided program synthesis. *Proc. ACM Program. Lang.* 6, OOPSLA1, Article 71 (April 2022), 27 pages. doi:10.1145/3527315
- [27] Pablo Antonio Martínez, Jackson Woodruff, Jordi Armengol-Estapé, Gregorio Bernabé, José Manuel García, and Michael F. P. O’Boyle. 2023. Matching Linear Algebra and Tensor Code to Specialized Hardware Accelerators. In *Proceedings of the 32nd ACM SIGPLAN International Conference on Compiler Construction* (Montréal, QC, Canada) (CC 2023). Association for Computing Machinery, New York, NY, USA, 85–97. doi:10.1145/3578360.3580262
- [28] Daye Nam, Baishakhi Ray, Seohyun Kim, Xianshan Qu, and Satish Chandra. 2022. Predictive synthesis of API-centric code. In *Proceedings of the 6th ACM SIGPLAN International Symposium on Machine Programming* (San Diego, CA, USA) (MAPS 2022). Association for Computing Machinery, New York, NY, USA, 40–49. doi:10.1145/3520312.3534866
- [29] Maxwell Nye, Luke Hewitt, Joshua Tenenbaum, and Armando Solar-Lezama. 2019. Learning to Infer Program Sketches. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, 4861–4870. <https://proceedings.mlr.press/v97/nye19a.html>
- [30] M.F.P. O’Boyle and P.M.W. Knijnenburg. 2002. Integrating Loop and Data Transformations for Global Optimization. *J. Parallel Distrib. Comput.* 62, 4 (April 2002), 563–590. doi:10.1006/jpdc.2001.1815
- [31] Augustus Odena, Kensen Shi, David Bieber, Rishabh Singh, Charles Sutton, and Hanjun Dai. 2021. BUSTLE: Bottom-Up Program Synthesis Through Learning-Guided Exploration. arXiv:2007.14381 [cs.PL] <https://arxiv.org/abs/2007.14381>
- [32] Theo X. Olausson, Jeevana Priya Inala, Chenglong Wang, Jianfeng Gao, and Armando Solar-Lezama. 2024. Is Self-Repair a Silver Bullet for Code Generation?. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=y0GJXRungR>
- [33] OpenAI. 2024. GPT-4 Technical Report. arXiv:2303.08774 [cs.CL] <https://arxiv.org/abs/2303.08774>
- [34] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems*, Vol. 32. Curran Associates, Inc., 8026–8037. <https://dl.acm.org/doi/10.5555/3454287.3455008>
- [35] Jie Qiu, Colin Cai, Sahil Bhatia, Niranjana Hasabnis, Sanjit A. Seshia, and Alvin Cheung. 2024. Tenspiler: A Verified-Lifting-Based Compiler for Tensor Operations. In *38th European Conference on Object-Oriented Programming (ECOOP 2024) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 313)*, Jonathan Aldrich and Guido Salvaneschi (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 32:1–32:28. doi:10.4230/LIPIcs.ECOOP.2024.32
- [36] Jonathan Ragan-Kelley, Connelly Barnes, Andrew Adams, Sylvain Paris, Frédo Durand, and Saman Amarasinghe. 2013. Halide: a language and compiler for optimizing parallelism, locality, and recomputation in image processing pipelines. *Acm Sigplan Notices* 48, 6 (2013), 519–530. doi:10.1145/2491956.2462176
- [37] Kensen Shi, David Bieber, and Rishabh Singh. 2022. Tf-coder: Program synthesis for tensor manipulations. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 44, 2 (2022), 1–36. doi:10.1145/3517034
- [38] Kensen Shi, Hanjun Dai, Kevin Ellis, and Charles Sutton. 2022. CrossBeam: Learning to Search in Bottom-Up Program Synthesis. arXiv:2203.10452 [cs.LG] <https://arxiv.org/abs/2203.10452>

Received 2024-11-15; accepted 2025-03-06