

LLASP: Fine-tuning Large Language Models for Answer Set Programming

Erica Coppolillo^{1,2}, Francesco Calimeri^{1,3}, Giuseppe Manco², Simona Perri¹, Francesco Ricca¹

¹University of Calabria, Via Bucci 28c, Rende - Italy

²ICAR-CNR, Via Bucci 8/9C, Rende - Italy

³DLVSystem Srl, Viale della Resistenza 19C, Rende - Italy

{erica.coppolillo, francesco.calimeri, simona.perri, francesco.ricca}@unical.it,
giuseppe.manco@icar.cnr.it

Abstract

Recently, Large Language Models (LLMs) have showcased their potential in various natural language processing tasks, including code generation. However, while significant progress has been made in adapting LLMs to generate code for several imperative programming languages and tasks, there remains a notable gap in their application to declarative formalisms, such as Answer Set Programming (ASP). In this paper, we move a step towards exploring the capabilities of LLMs for ASP code generation. First, we perform a systematic evaluation of several state-of-the-art LLMs. Despite their power in terms of number of parameters, training data and computational resources, empirical results demonstrate inadequate performance in generating correct ASP programs. Therefore, we propose LLASP, a fine-tuned lightweight model specifically trained to encode fundamental ASP program patterns. To this aim, we create an ad-hoc dataset covering a wide variety of fundamental problem specifications that can be encoded in ASP. Our experiments demonstrate that the quality of ASP programs generated by LLASP is remarkable. This holds true not only when compared to the non-fine-tuned counterpart but also when compared to the majority of eager LLM candidates, particularly from a semantic perspective. All the code and data used to perform the experiments are publicly available: <https://github.com/EricaCoppolillo/LLASP>.

1 Introduction

Answer Set Programming (ASP) (Brewka, Eiter, and Truszczyński 2011; Lifschitz 2019) is a pivotal tool for knowledge representation and reasoning, offering a range of benefits that contribute to its significance in various domains. One notable advantage is ASP’s capacity to handle knowledge in a concise and intuitive manner, allowing users to express complex problems with relative ease (Baral 2010). ASP provides a declarative approach for expressing knowledge and solving combinatorial optimization, planning, and reasoning tasks (Gebser et al. 2012). Its non-monotonic nature allows for the representation of incomplete information and default reasoning, essential for capturing real-world scenarios. Moreover, ASP’s expressiveness enables the integration of various types of knowledge, including rules, constraints, and preferences, facilitating flexible problem-solving (Lifschitz 2019).

This formalism is typically regarded as a specialized tool

utilized by domain experts and knowledge engineers familiar with specific problem domains. These individuals leverage ASP as a declarative language to articulate problem constraints, rules, and preferences succinctly and intuitively. However, despite the language’s simplicity and conciseness, coding in ASP can be an overwhelming obstacle for non-expert users, due to lack of familiarity with declarative paradigms and their intrinsic semantics, as well as potential complexity of application domains. Under this perspective, it is important to develop tools that enhance the efficiency and automation of creating ASP programs. These tools would minimize the disparity between natural language specifications and the corresponding ASP source code, facilitating smoother development processes. (Erdem and Yeniterzi 2009; Fang and Tompits 2017; Schwitter 2018; Caruso et al. 2024).

In this respect, the recent advancements in Artificial Intelligence and Machine/Deep Learning have led to the emergence of Large Language Models (LLMs) as indispensable assets across various natural language applications and tasks (Raiaan et al. 2024; Minaee et al. 2024). Notably, in the domain of automatic text-to-code translation, LLMs are employed to generate programs from natural language inputs. These models are trained on extensive datasets encompassing code repositories, technical forums, coding platforms, documentation, and relevant web data associated with programming. Such comprehensive training equips LLMs with the ability to grasp the nuances of code context, thereby enhancing the accuracy of contextually relevant code generation.

Indeed, a natural prosecution in this context is the investigation of LLMs-based approaches aimed at automating code generation within the context of declarative programming. This research challenge aims at narrowing the gap between natural language specifications and ASP source code, and aligns with the broader trend of leveraging language models for code automation. Further, leveraging LLMs for automating ASP program generation offers several advantages, which include: better semantic specification through high-level natural language descriptions of their problem domains; contextual awareness and interoperability, which simplifies the process of specifying ASP programs, making it accessible to a wider range of users without deep programming expertise; and enhanced efficiency

and productivity. Surprisingly, the current literature exhibits a significant discrepancy in this respect, even if some initial attempts have been made (Ishay, Yang, and Lee 2023; Borroto, Kareem, and Ricca 2024).

In an effort to rectify this gap, the present work proposes a systematic study on the capabilities of currently available LLMs to automate the generation of ASP programs. In our study, we identify some structural properties (classes) of ASP programs specifically tailored to fundamental tasks, and correlate them with the generation capabilities of the currently available LLM architectures. Inspired by such a characterization, we further devise a methodology for obtaining more accurate ASP encodings. Our approach consists in the construction of a custom training dataset, specifically tailored for the identified relevant tasks, and in fine-tuning an LLM instance that can be successfully exploited for ASP coding. Through an extensive experimentation, we show that even a fine-tuned lightweight base-model is more effective than heavy-sized LLMs. As a result, the proposed methodology offers a pathway towards fine-tuned models ready for practical deployment.

Our contributions can be summarized as follows.

- To the best of our knowledge, we are the first to extensively compare state-of-the-art LLMs in terms of ASP code generation. Our findings suggest that, despite their potential, LLMs still underperform in terms of syntactic and semantic program correctness.
- We show that a tailored training strategy, even if applied on a lightweight LLM, can outperform state-of-the-art large-size models. To address this, we curated a comprehensive dataset and used it to produce LLASP, the fine-tuned version of a Gemma 2B base-model specifically trained to capture fundamental ASP patterns.
- Via a comprehensive experimental evaluation, we prove that LLASP generates ASP programs that are highly valuable in terms of both syntactic and semantic accuracy, overcoming significantly larger and more powerful LLMs especially under a semantic perspective.
- Finally, we provide a discussion on the current limitations and potential directions for extensions of our proposed approach, inspired by the results of additional experiments.

The rest of the paper is structured as follows. Section 2 discusses the state of the art in automated code generation and the drawbacks of current approaches focused on ASP. In Section 3, we introduce the basic concepts, and in Section 4 we propose our methodology and design patterns for the evaluation and fine-tuning of the model. Section 5 reports the experiments and discusses our findings. Finally, in Section 6 we set some pointers for future research.

2 Related Work

The advantages of automating code synthesis are well-recognized in the literature (Ernst and Bavota 2022; Kalliamvakou 2022; Peng et al. 2023; Dakhel et al. 2023), and nearly all mainstream programming languages are nowadays supported by automated program composition tools (Chen, Tworek, and et al. 2021). In this area, Large

Language Models (LLMs) play a central role, and their performance has been extensively compared in the literature (Xu et al. 2022; Wang et al. 2023). The positive impact of fine-tuning models for code generation of imperative programming is also established (Ma et al. 2024).

In the context of declarative programming, a widely recognized research objective is to create tools that streamline and automate the development of Answer Set Programming (ASP) programs. The goal is to bridge the gap between natural language specifications and ASP source code (Erdem and Yeniterzi 2009; Fang and Tompits 2017; Schwitter 2018; Caruso et al. 2024). Initial proposals focused on automating the resolution of logic puzzles presented in simplified English by translating their descriptions into ASP (Baral and Dzifcak 2011), employing λ -calculus and probabilistic combinatorial categorical grammars. Later, several efforts have been spent in the development of Controlled Natural Languages (CNLs) (Kuhn 2014) for ASP programs. CNLs represent subsets of full natural languages, featuring restricted grammar and vocabulary. Among them, the BIOQUERYCNL (Erdem and Yeniterzi 2009) defines the grammatical structure of a CNL and algorithm for transpiling queries into ASP; Fang et al. introduced a CNL approach for ASP that leverages LANA annotations, that was implemented in the SeaLion IDE (Fang and Tompits 2017). In 2018, Schwitter developed a CNL called PENG^{ASP} for specifying and verbalizing answer set programs (Schwitter 2018). More recently, Dodaro et al. introduced CNL2ASP, an extensive publicly-available tool for converting controlled natural language into ASP programs (Caruso et al. 2024). On the one hand, the CNLs provide the programmer with a language that is more similar to the natural one, thus greatly reducing the barrier of coding in ASP code; on the other hand, CNLs still have their own syntactic constraints that do not completely free the user from structuring the sentence according to a formal syntax.

Even powerful language tools such as LLMs have been exploited in conjunction with the ASP formalism in several ways. (Nye et al. 2021) introduced a dual-system model based on GPT-3, comprising neural System 1 and logical System 2, which generates semantic parsers from natural language sentences and integrates them with reasoning modules. In a similar vein, (Yang, Ishay, and Lee 2023) proposed that LLMs like GPT-3 can function as few-shot semantic parsers, transforming natural language into logical forms for ASP without necessitating distinct retraining for diverse question-answering tasks. (Ishay, Yang, and Lee 2023) utilized LLMs with prompt engineering to obtain ASP solutions for logic puzzles, leveraging the logic puzzle dataset from (Mitra and Baral 2016). Further, the effective utilization of Large Language Models in combination with ASP to perform a number of natural language understanding tasks has been explored and implemented in the STAR framework (Rajasekharan et al. 2023).

Despite the growing interest in this area, however, a notable gap still persists in the literature regarding the adoption of LLMs for code synthesis for declarative programming languages, such as ASP. In 2024, Borroto et al. moved the first steps in this regard. Their method, implemented in the

NL2ASP tool, constructs ASP programs from natural language specifications through a two-step architecture (Boroto, Kareem, and Ricca 2024). Indeed, NL2ASP first applies neural machine translation (Stahlberg 2020) to transform natural language into the statements of the CNL by Dodaro et al.; then, in the second step, the produced sentences are used to generate ASP code by means of the CNL2ASP tool. NL2ASP has been implemented with two Transformer-based models for NMT tasks, i.e., T5-small, and Bart-base, and demonstrated promising performance.

Notably, there are some key differences between the approach proposed in this paper and the cited proposal. First, NL2ASP currently focuses on generating ASP programs for solving graph-related problem specifications, whereas the approach presented in this paper does not focus on a specific domain. Rather, it targets the generation of ASP programs suitable for solving general problems, starting from some fundamental specific ASP patterns, as we will show later. Another important difference is that our approach does not rely on an intermediate format: rather, it aims at producing ASP code directly from natural language specifications. Finally, it is worth observing that, even though low performance of LLMs in generating ASP code was already preliminarily observed (Boroto, Kareem, and Ricca 2024), a systematic and exhaustive empirical analysis assessing this behaviour was still missing, and is provided in this paper.

3 Preliminaries

Large Language Models. A Large Language Model \mathcal{L} can be formally described as a function f , which stochastically maps input sequences of tokens $x = [x_1, x_2, \dots, x_n]$ to an output sequence $y = [y_1, y_2, \dots, y_m]$, where n represents the length of the input sequence and m represents the length of the output sequence. The model defines $P_{\mathcal{L}}(y|x)$, i.e., the probability distribution of y given x , capturing complex patterns and relationships in natural language. Since $x \in V^*$ and $y \in W^*$, where V, W represent vocabularies of tokens, both x and y should be compliant with respect to some grammar. The f function is finally defined based on a sampling of y from $P_{\mathcal{L}}(\cdot|x)$.

The above probability is computed by relying on the Transformer Encoder-Decoder architecture introduced in (Vaswani et al. 2017). In more details, the input sequence of tokens is first converted into dense vector representations, which encode context similarity into geometrical closeness and also embed positional information. These embeddings then feed the Transformer layers. The Encoder consists of multiple identical layers, each with self-attention to weigh word importance, and a fully connected feedforward network for independent position transformations. Similarly, Decoder layers feature self-attention with masking, Encoder-Decoder Attention for input focus during token generation, and position-wise feedforward networks. Finally, the output of the Transformer Decoder is passed through a linear layer followed by a softmax activation function, which generates a probability distribution over the vocabulary of possible output tokens.

The latter implements the function $P_{\mathcal{L}}(y_i|x, y_{i-1})$ upon which the sampling process can be devised as a sequence of

steps. Here, y_{i-1} represents the prefix of y up to position $(i-1)$. Several architectural variants have been proposed, e.g., Causal-Decoder, Prefix Decoder, Autoregressive, Mixture Variants (Naveed et al. 2024; Raiaan et al. 2024). The underlying architecture of the LLM is not deeply investigated in this paper, since it is orthogonal to our study.

The training of \mathcal{L} is usually accomplished in steps. The first step consists in a pre-training task which is based on next-word prediction. In practice, given x and a partial response y_{i-1} , the task is to learn to correctly predict y_i . The next phase is called supervised fine-tuning (SFT), and it consists in refining \mathcal{L} to predict the whole y given x , on a large corpus $\mathcal{D} = \{(x^{(1)}, y^{(1)}), \dots, (x^{(N)}, y^{(N)})\}$, where N is the number of training pairs. The refinement is based on the loss $\ell(x^{(i)}, P_{\mathcal{L}}(y^{(i)}|x^{(i)}))$, usually defined in terms of cross-entropy:

$$\ell(y^{(i)}, P_{\mathcal{L}}(y^{(i)}|x^{(i)})) = \sum_j y_j^{(i)} \log P_{\mathcal{L}}(y_j^{(i)}|x^{(i)}, y_{j-1}^{(i)})$$

In the following, we shall refer to this formulation of the loss, even though alternative formulations are possible.

Answer Set Programming. ASP is a powerful declarative formalism for Knowledge Representation and Reasoning that gained increasing interest for its high expressive power, and the availability of solid and effective implementations (Gebser et al. 2018). It is fully declarative (i.e., the ordering of literals and rules is immaterial), and the encoding of a large variety of problems is simple and concise.

In the following, we briefly recall the syntax of the language, focusing on the aspects relevant to this work, and provide an intuitive semantics, revisiting the concept of *answer sets* that will play a crucial role in the validation process outlined in Section 5.2. For further details and complete references to advanced ASP features, the reader may refer to the vast literature (Brewka, Eiter, and Truszczyński 2011; Eiter, Ianni, and Krennwallner 2009; Calimeri et al. 2016).

A *term* is either a constant or a variable. An *atom* is an expression $p(t_1, \dots, t_n)$, where p is a predicate of arity n and t_1, \dots, t_n are terms. A program P is a finite set of rules, constructs like *Head* : - *Body* where *Head* is a disjunction of atoms and *Body* is a conjunction of literals; more formally, a rule is of the form

$$\alpha_1 | \dots | \alpha_k : - \beta_1, \dots, \beta_n, \text{not } \beta_{n+1}, \dots, \text{not } \beta_m.$$

where $m \geq 0, k \geq 0; \alpha_1, \dots, \alpha_k$ and β_1, \dots, β_m are atoms. A rule is interpreted according to common sense principles: roughly, its intuitive semantics corresponds to an implication. When a predicate p occurs in the head of a rule r , we say that r defines p . Rules featuring an atomic formula in the head and an empty body are used to represent information known to be certainly true and indeed called *facts*. A rule with empty head is called (*strong*) *constraint*, and intuitively expresses conditions that must be satisfied; basically, it is forbidden that all the literals in the body of a constraint are true. ASP also supports *weak constraints*, i.e., special rules with empty heads which allow to express preferences, and hence to deal with optimization problems. A weak constraint of the form

$$:\sim l_1, \dots, l_n. [w@l, t_1, \dots, t_m]$$

associates literals l_1, \dots, l_n with a weight w , a level l , and additional terms t_1, \dots, t_m for $m \geq 0$. Intuitively, weak constraints allow to express conditions that should be satisfied, but not necessarily have to be; if a weak constraint is violated (i.e., its body is true), then the weight w (“cost”) is paid at priority level l . Finally, a program (a rule, an atom) that does not contain variables is said to be *ground*.

Models are defined over ground programs: a *model* for a program P is a subset of all possible ground atoms that satisfies all rules in P . The semantics of ASP programs is given in terms of special models called *Answer Sets*: according to this semantics, an ASP program may have several alternative answer sets (but possibly none), each corresponding to a possible view of the world (Gelfond and Lifschitz 1991; Brewka, Eiter, and Truszczyński 2011). In ASP, a computational problem is typically solved by modeling it via a program consisting of a collection of rules along with a set of facts representing the instance at hand; then, solutions are found by computing the intended answer sets, according to the so-called answer set semantics. Answer sets correspond one-to-one to the solutions of the given instance of the modeled problem: if a program has no answer sets, the corresponding instance has no solutions. If weak constraints are present, optimal answer sets are those minimizing the sum of weights of the violated weak constraints in the highest priority level. Among these, the optimal sets are those that also minimize the sum of weights of the violated weak constraints at the next lower level, and so forth.

4 Methodology and Knowledge Design

In the present work, we address a twofold objective. First, we focus on pre-trained general-purpose Large Language Models, not specifically tailored for ASP generation. Thus, given an LLM \mathcal{L} and a natural language specification x of a problem, we aim at evaluating the robustness of \mathcal{L} in sampling $y \sim P_{\mathcal{L}}(\cdot|x)$, such that y is an ASP encoding compliant to x . Our second objective revolves around exploring how fine-tuning impacts the quality of generation, particularly when dealing with small-sized models. For a comprehensive evaluation, we need to consider various aspects:

- **Dataset diversity and complexity.** We need to assemble a dataset encompassing a wide range of problem descriptions (x) covering different domains, complexities, and lengths. This diversity ensures that the evaluation captures \mathcal{L} ’s performance across various contexts and complexity levels, thus helping in discerning \mathcal{L} ’s efficacy in handling various ASP encoding challenges.
- **ASP compliance and robustness.** Since each problem description x will be paired with a corresponding ASP encoding $y \sim P_{\mathcal{L}}(\cdot|x)$, we need to verify whether y adheres to the ASP syntax, as well as its fidelity to the original problem description x . Metrics such as syntax accuracy, semantic coherence, and adherence to problem constraints should be considered.
- **Impact of model size.** We need to investigate how the size and architecture of \mathcal{L} influence its ability to generate accurate ASP encodings. This analysis should involve comparing the performance of different \mathcal{L} variants

(in particular, smaller vs. larger models) and assessing any trade-offs between model size and encoding quality.

- **Generalizability assessment.** We need to examine whether \mathcal{L} , trained on a diverse set of tasks, can generalize well to new problem domains not encountered during training. This helps devise the extent to which \mathcal{L} ’s capabilities extend beyond its training data and whether it can effectively tackle novel ASP encoding tasks.

We concentrate on the first aspect for now, addressing the others in the subsequent section. Here, we need to define an ad-hoc approach for tuning \mathcal{L} towards the production of ASP programs. In this regard, we take inspiration from the typical learning process experienced by humans while trained on the task of encoding problems into ASP programs. In particular, the process is based on modeling simple patterns encoding basic pieces of knowledge. For instance, how to represent with ASP a Cartesian product, the join between the extensions of two predicates, simple guesses, conditions on specific values, transitive closures, and so on. In principle, more complex patterns can be devised in a compositional way, starting from these basic building blocks to produce larger sets of rules that, in turn, encode more complex knowledge. This approach to ASP generation is in general effective: the declarative nature of the formalism makes the order between rules in a program, and the order of literals in the bodies, immaterial; thus, in many cases, a program representing the desired solution can be composed by just adding subprograms representing additional knowledge.

According to this idea, we defined a template set $\mathcal{T} = (\mathcal{P}, \mathcal{A})$, where \mathcal{P} is a set of descriptions of basic tasks expressed in natural language and \mathcal{A} is a set of text snippets consisting of corresponding (*gold*) ASP programs that encode them. In practice, \mathcal{T} is partitioned into $\{\mathcal{T}_{C^1}, \mathcal{T}_{C^2}, \dots, \mathcal{T}_{C^k}\}$, where C^i , represents the i -th task type and k is the total number of modeled task types. For instance, if C^1 refers to a transitive closure task, \mathcal{T}_{C^1} consists of pairs $(\mathcal{P}_{C^1}, \mathcal{A}_{C^1})$ featuring a description in natural language and one among the corresponding suitable gold ASP encodings, respectively. Each template (and its gold ASP counterpart) exhibits some placeholders to be instantiated. Such placeholders represent predicates, labels and values within the ASP encoding.

Below, we provide a brief overview of the fundamental tasks we examined. For each task, we offer an informal description, the template utilized, and an example of instantiated prompt along with its corresponding ASP encoding.

Guessing Assignments. We started with considering a typical task where all elements of a given set must be assigned to a unique label picked from a fixed set. A common way for representing this in ASP is to use disjunction in order to express a so-called “guess”. One possible pair associating a prompt with a suitable ASP encoding is reported next.

```

Template: Write an ASP program for the following
problem. Assign exactly a label among a given set
of labels to a set of elements. The set of
elements is expressed by predicate [PREDICATE].
The labels are [LABEL]+.

```

Prompt: Write an ASP program for the following problem. Assign exactly a label among a given set of labels to a set of elements. The set of elements is expressed by predicate `city`. The labels are `moscow,rome,dubai`.

```
Encoding: assign(X,"moscow") | assign(X,"rome") |
         assign(X,"dubai") :- city(X).
```

Expressing Constraints. Another typical task is to express conditions that must be fulfilled. Within ASP, this is commonly expressed by means of (classical/strong) constraints, possibly associated with auxiliary rules if needed. One possible pair associating a prompt asking for preventing a specific assignment with a suitable ASP encoding is reported next.

Template: Write an ASP program for the following problem. Prevent the predicate `[PREDICATE]` with value `[VALUE]` from having label `[LABEL]`.

Prompt: Write an ASP program for the following problem. Prevent the predicate `car` with value `11` from having label `"red"`.

```
Encoding: :- assign(11,"red").
```

Generating Combinations. The need for generating all combinations of elements from two different sets (i.e., defining the Cartesian product) can be in general fulfilled in ASP by means of a simple rule. One possible pair associating a prompt with a suitable ASP encoding is reported next.

Template: Write an ASP program for the following problem. Generate all the combinations of elements from two sets. The two sets are represented by predicates `[PREDICATE_1]` and `[PREDICATE_2]`.

Prompt: Write an ASP program for the following problem. Generate all the combinations of elements from two sets. The two sets are represented by predicates `city` and `airport`.

```
Encoding: combination(X,Y) :- city(X), airport(Y).
```

Joins. Joins over the elements of two different sets according to specific matching criteria over features can be typically defined in ASP via rules featuring proper body literals having a variable in common. One possible pair associating a prompt with a suitable ASP encoding is the following.

Template: Write an ASP program for the following problem. Consider predicate `[PREDICATE_1]` having fields `[LABEL]+` and the predicate `[PREDICATE_2]` having fields `[LABEL]+`. Define a predicate `[PREDICATE_1]_[PREDICATE_2]` that associates to each `[PREDICATE_1]` the `[LABEL]` of `[PREDICATE_2]`.

Prompt: Write an ASP program for the following problem. Consider predicate `"owner"` having fields `"ID","surname","name","restaurantID"`, and the predicate `"restaurant"` having fields `"ID","description"`. Define a predicate `"owner_restaurant"` that associates to each owner the description of restaurant.

```
Encoding: owner_restaurant(X,Z) :- owner(X,_,_,Y),
                                             restaurant(Y,Z).
```

Transitive Closure. Transitive closure is a fundamental tool for defining the structure of relationships in a variety of contexts as it catches not only direct relationships among elements, but also those that are indirectly resulting from a chain of relations. Expressing this with ASP in general requires to use more than a single rule. One possible pair associating a prompt describing a transitive closure with a suitable ASP encoding is reported next; in this case, the encoding make use of two rules: one for defining the direct relations, and another (relying on recursion) for catching the indirect ones.

Template: Write an ASP program for the following problem. Define predicate `[PREDICATE_1]` as the transitive closure of predicate `[PREDICATE_2]`.

Prompt: Write an ASP program for the following problem. Define predicate `"arrivals"` as the transitive closure of predicate `"flight"`.

```
Encoding:
         arrivals(X,Y) :- flight(X,Y).
         arrivals(X,Y) :- flight(X,Z),arrivals(Z,Y).
```

Expressing Preferences. In order to use ASP for solving optimization problems, one must be able to express, to several extents, preferences over the set of admissible solutions; this is usually done by encoding programs featuring *weak constraints*. One possible pair associating a prompt asking for expressing a specific preference with a suitable ASP encoding is reported next.

Template: Write an ASP program for the following problem. I would prefer that predicate `[PREDICATE]` with value `18` is not associated with `[LABEL]`. If this occurs, it costs `[COST]` at level `[LEVEL]`.

Prompt: Write an ASP program for the following problem. I would prefer that predicate `house` with value `18` is not associated with `"flat"`. If this occurs, it costs `2` at level `2`.

```
Encoding: :~ assign(18,"flat"). [2@2]
```

Filtering. When encoding ASP programs, it is frequently necessary to apply filters to the extensions of certain predicates based on specific requirements. We report next some common types of filtering criteria.

Filtering by values. The first type of filtering criteria consists in selecting the portion of the extension of a predicate that match with a specific value. One possible pair associating a prompt with a suitable ASP encoding is reported next.

Template: Write an ASP program for the following problem. Select all values associated to the predicate `[PREDICATE]` with label `[LABEL]`.

Prompt: Write an ASP program for the following problem. Select all values associated to the predicate `color` with label `"purple"`.

```
Encoding: select(X) :- color(X, "purple").
```

Filtering by negative. Another type of filtering criteria consists in excluding the portion of the extension of a predicate that matches with a given condition. This could consist, for instance, of a difference (subtraction) between two predicate extensions, but also the negation of compound conditions. One possible pair associating a prompt asking for filtering a predicate according to a negative condition defined over a filtered portion of a different predicate with a suitable ASP encoding is reported next.

```
Template: Write an ASP program for the following
problem. Select all values associated with
predicate [PREDICATE_1] but not associated with
predicate [PREDICATE_2] and label [LABEL].
```

```
Prompt: Write an ASP program for the following
problem. Select all values associated with
predicate vehicle but not associated with
predicate moto and label "kawasaki".
```

```
Encoding: select(X) :- vehicle(X),
not moto(X, "kawasaki").
```

Filtering by numeric comparisons. Filtering portion of tables according to the results of comparisons between terms is another typical task. One possible pair associating a prompt asking for filtering a predicate according to a numeric comparison with a suitable ASP encoding is reported next.

```
Template: Write an ASP program for the following
problem. Select all values associated with
predicate [PREDICATE] with a value greater or
equal than [VALUE].
```

```
Prompt: Write an ASP program for the following problem.
Select all values associated with predicate
size with a value greater or equal than 5.
```

```
Encoding: select(X) :- size(X,C), C>=5.
```

This is a small yet rich set of fundamental pieces of knowledge that an ASP user can draw upon and combine so to build more complex programs for solving different problems. As an example, think of many combinatorial problems, that basically consist in selecting or associating elements from a given collection in ways that comply with a given set of constraints; suitable programs for such problems are easily obtainable by combining *guessing assignments* and *expressing constraints*, possibly relying on the other types of tasks introduced above. If also optimization plays a role, suitable solutions can be obtained by additionally *expressing preferences*.

5 Experiments

In this section, we design a suite of experiments and evaluate their results. For the prompts described in the previous section, we aim at evaluating the following.

- **RQ1:** To what extent are state-of-the-art pre-trained LLMs actually capable of encoding textual problems into accurate ASP programs?

- **RQ2:** Can LLMs be fine-tuned for generating ASP rules that comply to a textual description?
- **RQ3:** To what extent can the output of the fine-tuned LLM be considered “ready-to-use” for practical applications, in terms of syntactic and semantic correctness?
- **RQ4:** What are the limitations and potentials of this approach, e.g., how do the prompt structure and/or the complexity of the problem affect the generated programs?

In order to ensure reproducibility, all the code and data used to perform the experiments are publicly available: <https://anonymous.4open.science/r/LLASP-D86C/>.

5.1 Pre-trained LLMs

Our initial goal is to assess the performance of several LLMs in generating ASP programs. Thus, in our experiments, we provide a structured evaluation of their performance in this domain. We focus on freely available models that also exemplify the latest advancements in LLM-based generation:

- **ChatGPT 3.5**¹: released by OpenAI, it is a fine-tuned version of GPT-3.5, a language model trained to produce text, that has been optimized for conversation.
- **Copilot**²: it is a conversational chat interface developed by Microsoft, built upon the language model GPT-4 and the text-to-image model DALL-E 3.
- **Gemini** (Team 2024a)³: a new class of multimodal LLMs introduced by Google, showing remarkable capabilities across image, audio, video, and text understanding.
- **Gemma** (Team 2024b)⁴: released by the same group that introduced Gemini, Gemma is a lightweight family of LLMs that outperforms similarly sized models across different tasks, such as understanding, reasoning and safety.
- **LLaMa2** (Touvron and others 2023)⁵: an LLM optimized for dialogue applications, developed by Meta. It shows competitive results across various benchmarks, consistently outperforming open-source chat models.
- **LLaMa3**⁶: an enhanced state-of-the-art version of LLaMa models, showing improvements in tasks like reasoning, code generation and adherence to instructions.
- **Mistral** (Jiang et al. 2023)⁷: despite the reduced size of the first release (7B parameters), Mistral is a LLM that showed superior performance and efficiency across several benchmark tasks. Since then, bigger and more powerful versions have been introduced.

Table 1 briefly compares the chosen LLMs in terms of model size, architectural details and training sources (Minae et al. 2024). We can identify a group of large size models (≥ 70 B parameters) and another group of reduced

¹<https://chatgpt.com/>

²<https://learn.microsoft.com/en-us/copilot/overview>

³<https://gemini.google.com/app>

⁴<https://huggingface.co/blog/gemma>

⁵<https://www.llama2.ai/>

⁶<https://llama.meta.com/llama3/>

⁷<https://chat.mistral.ai/chat>

size models (< 70B parameters); Gemma 2B is the smallest model in the batch.

Model	No. Params.	Architecture	Training Data
ChatGPT 3.5	175B	E-D	Online sources
Copilot	1.5T	E-D	Github repositories
Gemini	1.6T	MoE	Docs, Books, Code
Gemma	2B-7B	D-Only	Docs, Maths, Code
LLaMa2	7B-13B	D-Only	Online sources
LLaMa3	8B-70B	D-Only + GQA	Online sources
Mistral	7B-141B	D-Only + GQA	Online sources

Table 1: Details on the evaluated pre-trained LLMs. Number of parameters is reported in billions(B)/trillions(T). E-D refers to the Encoder-Decoder architecture; D-Only is Decoder-Only; MoE is Mixture of Experts, and GQA is Grouped-Query Attention.

5.2 LLASP Fine-tuning

The next phase of the proposed methodology involves fine-tuning a lightweight model and assessing it on the problem categories outlined in Section 4. To achieve this, we build a training dataset by generating instances of the templates relative to such categories. Basically, for each problem class \mathcal{T}_{C^i} , we produce variations of that template to generate a collection of prompts along with corresponding suitable ASP instances. Templates are instantiated by using a predefined set of predicates and labels, as shown in Section 4. The re-

Problem	No. Tuples	Proportion (%)
Assignment	1,000,000	27
Constraint	500,000	13.5
Combination	100,000	2.7
Join	900,000	24.3
Transitive closure	100,000	2.7
Preference	400,000	10.8
Value filtering	100,000	2.7
Negative filtering	100,000	2.7
Numeric filtering	500,000	13.5
Total	3,700,000	100

Table 2: Number of tuples and relative proportions of each problem within the dataset \mathcal{D} .

sulting dataset \mathcal{D} contains approximately 4 million tuples, with varying proportions based on the problem type, as detailed in Table 2. The quantity of tuples per task depends on its syntactic complexity: tasks supporting greater variability in their corresponding ASP programs (such as the number of predicates in the rule, arity of the atoms, and ground instances) require more data during training. We further split \mathcal{D} into a training and a validation set with a 80–20 ratio, maintaining the proportions per problem.

We opted for Gemma 2B as our model of choice for fine-tuning. As the smallest among the available models, we anticipate its baseline performance to be relatively modest; this allows us to obtain tangible feedback on the effectiveness of the methodology. We conduct Supervised Fine-Tuning (SFT) of \mathcal{L} using the `SFTTrainer` component from the

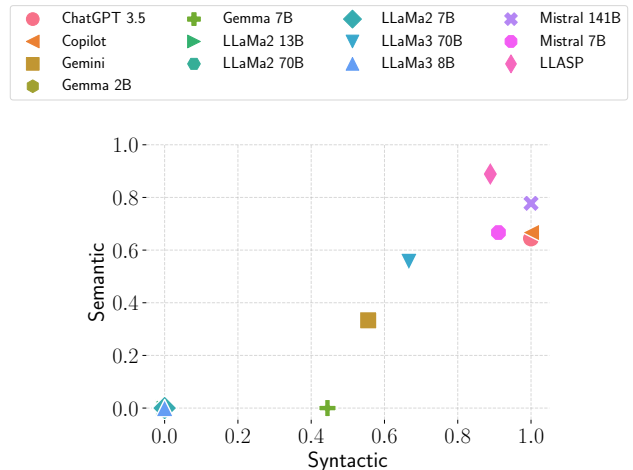


Figure 1: Comparative analysis on both syntactic and semantic results. Some markers are hidden for overlapping in (0, 0).

`trl` library⁸, while leveraging QLoRA (Dettrmers et al. 2023) for computational efficiency. The machine for training is an NVIDIA DGX server featuring a 20-Core Intel Xeon CPU, 256GB RAM and 4X NVIDIA Tesla V100 32 GB/GPU.

Evaluation Protocol. To assess the correctness of the generated ASP encodings, we rely on the ASP solver *Clingo* (Gebser et al. 2016) via the dedicated Python API⁹. We use *Clingo* for computing the answer sets of a program P , hence defining a function $s(P) = AS(P)$, where $AS(P)$ is the set (possibly empty) of answer sets of P . Given both the ASP program $y \sim P_{\mathcal{L}}(\cdot|x)$ generated by \mathcal{L} from the prompt x and the gold program y^* associated with x , we first build a set of facts F_{y^*} representing an instance of the textual problem x . Given $P = y \cup F_{y^*}$ and $P^* = y^* \cup F_{y^*}$, we then perform the following checks. We invoke *Clingo* for computing $s(P)$; if no parsing error occurs, we have a *syntactic* hit on y . Further, we compute $s(P)$ and $s(P^*)$, comparing $AS(P)$ with $AS(P^*)$; if they match, we have a *semantic* hit on y .

5.3 Results

In an initial series of experiments, we conduct a comparative analysis between the pre-trained models and LLASP (i.e., the fine-tuned Gemma 2B model). It is noteworthy that many pre-trained models do not offer direct invocation through APIs. Hence, the evaluation is conducted by selecting, for each task, the example instantiation outlined in Section 4. Then, the corresponding prompt is submitted, and results are assessed according to the aforementioned protocol. The process is repeated 5 times to ensure statistical robustness over the sampling process intrinsic of the LLMs.

Table 3 reports accuracy of all the considered models, considering both syntactic and semantic perspectives; Figure 1 offers a concise visual representation of these findings.

⁸<https://huggingface.co/docs/trl>

⁹<https://potassco.org/clingo/python-api/5.4/>

Model	Syntactic	Semantic
ChatGPT 3.5	1.	0.64
Copilot	1.	0.67
Gemini	0.56	0.33
Gemma 2B	0.	0.
Gemma 7B	0.44	0.
LLaMa2 7B	0.	0.
LLaMa2 13B	0.	0.
LLaMa2 70B	0.	0.
LLaMa3 8B	0.	0.
LLaMa3 70B	0.67	0.56
Mistral 7B	0.91	0.67
Mistral 141B	1.	0.78
LLASP	0.89	0.89

Table 3: Comparative analysis in syntactic and semantic terms.

Notably, a syntactic hit does not necessarily imply with semantic correctness. For instance, Gemma 7B achieves 44% syntactic accuracy but shows no semantic quality at all. In general, no model achieves complete correctness across the board, and lightweight models report the weakest performance (with the only notable exception of Mistral 7B).

Among the models, ChatGPT, Copilot, and Mistral 141B achieve 100% syntactic accuracy. Surprisingly, both Gemini and LLaMa3 70B display lower accuracy with respect to these methods, despite being comparable in model size. LLASP emerges as the model with the highest semantic accuracy among all the models considered. We note that whenever a generation produces a syntactically correct program, it also tends to be semantically correct. Table 3 confirms this observation, indicating that this phenomenon occurs only with LLASP. This is particularly noteworthy, given that the original Gemma 2B yields entirely inconsistent results. Table 4 further details the results for each task. Compared to large sized models, LLASP only fails on the join task, due to syntactic fails. By contrast, Mistral 141B exhibits semantic failures both on assignment and negative filtering.

We provide more insights regarding the generation flaws via anecdotal evidence of the ASP encoding provided by the identified top-performing models (see Figure 1): Mistral 141B, ChatGPT 3.5, Copilot, and Mistral 7B. In particular, we focus on semantic correctness of the generated programs, which in our context is the most representative quality measure. Surprisingly, almost all of the aforesaid model, except for ChatGPT 3.5, always fail in representing the guessing assignment via a semantically correct ASP encoding, thus producing a relative score of 0. (see Table 4). ChatGPT 3.5, indeed generated a semantically correct program 4 times out of 5, achieving a score of 0.8. We report below the encoding produced by each of the models, comparing it with the one generated by LLASP, in the response to the Guessing Assignment task reported in Section 4.

Mistral 7B/141B. Despite being syntactically correct, the generated program presents a semantic flaw, since it does not allow multiple assignments for the same label. This property is not requested in the prompt, and would lead to inconsistency whenever the number of labels is lower than the number of predicates.

```
Encoding:
1 { assigned(X, L) : label(L) } 1 :- city(X).
:- assigned(X1, L), assigned(X2, L), X1 != X2.
```

ChatGPT 3.5. Similarly, the result of the generation is a syntactically correct ASP encoding, according to the Clingo solver. However, it fails in capturing the semantics of the textual prompt.

```
Encoding:
1 { assign(X, L) : label(L) } 1 :- city(X).
:- assign(X, L1), assign(Y, L1), X != Y, label(L1).
:- city(X), not assign(X, _).
```

Copilot. Here, the generated program encodes a completely different semantics with respect to the problem description: indeed, the above fragment specifies a join between the predicate “city” and the predicate “label” via the “City” label.

```
Encoding:
assign_label(City, Label) :-
city(City), label(City, Label).
```

LLASP. Eventually, we report an example of wrong ASP program produced by LLASP over the join task discussed in Section 4. The generated encoding presents a syntactic error regarding both the arity of the predicate “owner” and the presence of the “Z” predicate, thus affecting the semantics of the program.

```
Encoding: owner_restaurant(X, Z) :- owner(X, Y), Z(Y).
```

As an additional test, in order to further assess the robustness of the findings, we compare LLASP against the aforesaid models by performing a more detailed evaluation on an extended set of test instances, where prompts vary in terms of predicates and labels. Table 5 reports the results. Again, despite occasional syntactic flaws, LLASP offers more reliable encodings, outperforming the second-best model (Mistral 7B) of almost ten percentage points. Moreover, we observe that the problem semantics is fulfilled whenever syntactic correctness is ensured.

Encouraged by these results, we proceeded in testing LLASP on a larger scale. To this purpose, we generated a test set \mathcal{D}_{test} consisting of 9,000 prompts, following the same principles as the training set \mathcal{D} described earlier, yet, again, on a different set of predicates and labels. In this test set, the number of examples for each task is evenly distributed. Table 6 presents the accuracy of LLASP over \mathcal{D}_{test} . We notice that: (i) unlike in the previous experiment, the consistency between syntactic and semantic accuracy is not upheld; (ii) the quality of the generated programs remains consistently high over all the tasks, even on the join problem which failed in our first experiment. Given the substantial scale of \mathcal{D}_{test} , we are confident that this experiment demonstrates the resilience of the proposed approach.

5.4 Limitations

To explore the ability of generating complex ASP programs from atomic tasks, we conducted further experiments

Model	Assignment		Constraint		Combination		Join		Closure		Preference		Value Filter		Neg. Filter		Num. Filter	
	Syn.	Sem.	Syn.	Sem.	Syn.	Sem.	Syn.	Sem.	Syn.	Sem.	Syn.	Sem.	Syn.	Sem.	Syn.	Sem.	Syn.	Sem.
ChatGPT 3.5	1.	0.8	1.	1.	1.	1.	1.	1.	1.	1.	1.	0.	1.	0.	1.	0.	1.	1.
Copilot	1.	0.	1.	1.	1.	1.	1.	1.	1.	1.	1.	0.	1.	1.	1.	0.	1.	1.
Gemini	0.	0.	1.	0.	0.	0.	1.	1.	1.	1.	0.	0.	0.6	0.	0.4	0.	1.	1.
Gemma 2B	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
Gemma 7B	0.	0.	1.	0.	0.	0.	1.	0.	0.	0.	0.	0.	0.	0.	1.	0.	1.	0.
LLaMa2 7B	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
LLaMa2 13B	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
LLaMa2 70B	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
LLaMa3 8B	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
LLaMa3 70B	0.8	0.2	0.4	0.4	1.	1.	0.6	0.6	1.	1.	0.	0.	0.6	0.6	0.6	0.2	1.	1.
Mistral 7B	0.8	0.	1.	0.	1.	1.	1.	1.	1.	1.	0.4	0.	1.	1.	1.	1.	1.	1.
Mistral 141B	1.	0.	1.	1.	1.	1.	1.	1.	1.	1.	1.	1.	1.	1.	1.	0.	1.	1.
LLASP	1.	1.	1.	1.	1.	1.	0.	0.	1.	1.	1.	1.	1.	1.	1.	1.	1.	1.

Table 4: Generation accuracy in terms of syntactic (*Syn.*) and semantic (*Sem.*) perspective across all the considered LLMs.

Model	Syntactic	Semantic
ChatGPT 3.5	1.	0.67
Copilot	1.	0.57
Mistral 7B	1.	<u>0.85</u>
Mistral 141B	1.	0.69
LLASP	<u>0.93</u>	0.93

Table 5: In-depth comparison using an extended dataset.

Problem	Syntactic	Semantic
Assignment	0.76	0.73
Constraint	1.	1.
Combination	1.	0.81
Join	0.95	0.91
Closure	1.	0.99
Preference	1.	0.88
Value Filtering	1.	0.89
Negative Filtering	1.	0.89
Numeric Filtering	1.	0.9
Total (avg)	0.97	0.89

Table 6: Results in terms of syntactic and semantic generation accuracy of LLASP over \mathcal{D}_{test} .

for highlighting limitations and potentials, by fine-tuning LLASP on an expanded version of the dataset \mathcal{D} , incorporating combined prompts. The results are mixed. While we observe a surprisingly successful generation rate for several complex combinations, we also witness failures on other simple ones. We believe that this is due to a suboptimal design of the training set regarding such combinations; indeed, we employed random combinations, whereas a more strategic approach should be devised.

We further assessed the generation’s generalization capabilities through a simple preliminary experiment: we generated and included 10 possible rephrasings for a given prompt in the dataset. However, the generation results were unsuccessful, suggesting that a more comprehensive evaluation requires a careful design of the training dataset.

Additionally, we explored various training strategies, in-

cluding Reinforcement Learning with Human Feedback (RLHF) (Ziegler et al. 2020). In this setup, generation is linked to a reward that considers both syntactic and semantic aspects of the ASP program generated. However, the results did not surpass those of standard SFT training. We posit that a more meticulous design of the reward function is needed, where several aspects of the generation (e.g., rule redundancy, fluency or verbosity) are to be considered.

6 Conclusions and Future Work

In this study, we conducted an extensive examination of the capability of Large Language Models (LLMs) to accurately encode Answer Set Programming (ASP) programs based on natural language specifications. To the best of our knowledge, this is the first extensive comparison of state-of-the-art LLMs in terms of ASP code generation. Our findings show that, despite the large potential, LLMs are not ready for such a task, as they underperform in terms of programs correctness (both syntactic and semantic). With this respect, we demonstrate that tailored training can significantly mitigate this deficiency, offering a pathway towards fine-tuned models able to generate ASP programs ready for practical deployment. Indeed, we propose LLASP, a fine-tuned version of the lightweight Gemma 2B base-model. We train it over a comprehensive dataset we designed and populated on purpose so to capture ASP fundamental patterns. LLASP-generated programs proved to be significantly better than those generated by larger and more powerful LLMs, especially under a semantic perspective.

We also devised and performed a significant set of additional experiments, that provide interesting insights on the topic and suggest that this study can be expanded in multiple directions. In particular, besides the pointers already discussed in Section 5.4 (i.e., generalizability, complex prompts, RLHF training), we claim that a novel metric is necessary, to be used both in training and evaluation. In principle, this metric should take into account both syntactic correctness and capability of the generated ASP programs to effectively tackle problems described by the prompts in the first place. Furthermore, we aim at developing novel controlled training methodologies to incorporate ASP syntac-

tic guidance within the training process of the LLMs. This would also require a careful inspection of the underlying architecture and the adoption of suitable modifications.

Finally, other valuable alternatives to fine-tuning warrant exploration, such as performing prompt engineering and training the model from scratch.

Acknowledgements

This work has been partially funded by: (1) MUR on D.M. 351/2022, PNRR Ricerca, CUP H23C22000440007; (2) PNRR project FAIR - Future AI Research (PE00000013), Spoke 9 - Green-aware AI, under the NRRP MUR program funded by the “NextGenerationEU”; (3) PNRR project Tech4You “Technologies for climate change adaptation and quality of life improvement”, CUP H23C22000370006, under the NRRP MUR program funded by the “NextGenerationEU”; PNRR Project “SoBigData.it – Strengthening the Italian RI for Social Mining and Big Data Analytics”, CUP B53C22001760006, under the NRRP MUR program funded by the “NextGenerationEU”; PNRR project SER-ICS - Security and Rights in the CyberSpace (PE00000014), CUP B29J24000680005, under the NRRP MUR program funded by the “NextGenerationEU”; PRIN PINPOINT CUP H23C22000280006. Francesco Calimeri is member of the Gruppo Nazionale Calcolo Scientifico-Istituto Nazionale di Alta Matematica (GNCS-INdAM).

References

- Baral, C., and Dzifcak, J. 2011. Solving puzzles described in english by automated translation to answer set programming and learning how to do that translation. In *AAAI Fall Symposium: Advances in Cognitive Systems*, volume FS-11-01 of *AAAI Technical Report*. AAAI.
- Baral, C. 2010. *Knowledge Representation, Reasoning and Declarative Problem Solving*. Cambridge University Press.
- Borroto, M.; Kareem, I.; and Ricca, F. 2024. Towards automatic composition of ASP programs from natural language specifications. In *IJCAI*, volume abs/2403.04541, to appear. [ijcai.org](https://www.ijcai.org).
- Brewka, G.; Eiter, T.; and Truszczynski, M. 2011. Answer set programming at a glance. *Commun. ACM* 54(12):92–103.
- Calimeri, F.; Gebser, M.; Maratea, M.; and Ricca, F. 2016. Design and results of the fifth answer set programming competition. *Artif. Intell.* 231:151–181.
- Caruso, S.; Dodaro, C.; Maratea, M.; Mochi, M.; and Riccio, F. 2024. CNL2ASP: converting controlled natural language sentences into ASP. *Theory Pract. Log. Program.* 24(2):196–226.
- Chen, M.; Tworek, J.; and et al., H. J. 2021. Evaluating large language models trained on code. *CoRR* abs/2107.03374.
- Dakhel, A. M.; Majdinasab, V.; Nikanjam, A.; Khomh, F.; Desmarais, M. C.; and Jiang, Z. M. J. 2023. Github copilot AI pair programmer: Asset or liability? *J. Syst. Softw.* 203:111734.
- Dettmers, T.; Pagnoni, A.; Holtzman, A.; and Zettlemoyer, L. 2023. Qlora: Efficient finetuning of quantized llms. Eiter, T.; Ianni, G.; and Krennwallner, T. 2009. Answer set programming: A primer. In Tesseract, S.; Franconi, E.; Eiter, T.; Gutierrez, C.; Handschuh, S.; Rousset, M.; and Schmidt, R. A., eds., *Reasoning Web. Semantic Technologies for Information Systems, 5th International Summer School 2009, Brixen-Bressanone, Italy, August 30 - September 4, 2009, Tutorial Lectures*, volume 5689 of *Lecture Notes in Computer Science*, 40–110. Springer.
- Erdem, E., and Yeniterzi, R. 2009. Transforming controlled natural language biomedical queries into answer set programs. In *BioNLP@HLT-NAACL*, 117–124. ACL.
- Ernst, N. A., and Bavota, G. 2022. Ai-driven development is here: Should you worry? *IEEE Softw.* 39(2):106–110.
- Fang, M., and Tompits, H. 2017. An approach for representing answer sets in natural language. In *DECLARE*, volume 10997 of *LNCS*, 115–131. Springer.
- Gebser, M.; Kaminski, R.; Kaufmann, B.; and Schaub, T. 2012. *Answer Set Solving in Practice*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.
- Gebser, M.; Kaminski, R.; Kaufmann, B.; Ostrowski, M.; Schaub, T.; and Wanko, P. 2016. Theory solving made easy with clingo 5. In *International Conference on Logic Programming*.
- Gebser, M.; Leone, N.; Maratea, M.; Perri, S.; Ricca, F.; and Schaub, T. 2018. Evaluation techniques and systems for answer set programming: a survey. In Lang, J., ed., *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, 5450–5456. [ijcai.org](https://www.ijcai.org).
- Gelfond, M., and Lifschitz, V. 1991. Classical negation in logic programs and disjunctive databases. *New Gener. Comput.* 9(3/4):365–386.
- Ishay, A.; Yang, Z.; and Lee, J. 2023. Leveraging large language models to generate answer set programs. In *KR*, 374–383.
- Jiang, A. Q.; Sablayrolles, A.; Mensch, A.; Bamford, C.; Chaplot, D. S.; de las Casas, D.; Bressand, F.; Lengyel, G.; Lample, G.; Saulnier, L.; Lavaud, L. R.; Lachaux, M.-A.; Stock, P.; Scao, T. L.; Lavril, T.; Wang, T.; Lacroix, T.; and Sayed, W. E. 2023. Mistral 7b.
- Kalliamvakou, E. 2022. Research: quantifying github copilot’s impact on developer productivity and happiness. *The GitHub Blog*.
- Kuhn, T. 2014. A survey and classification of controlled natural languages. *Comput. Linguistics* 40(1):121–170.
- Lifschitz, V. 2019. *Answer Set Programming*. Springer.
- Ma, Z.; Guo, H.; Chen, J.; Peng, G.; Cao, Z.; Ma, Y.; and Gong, Y.-J. 2024. Llamoco: Instruction tuning of large language models for optimization code generation.
- Minaee, S.; Mikolov, T.; Nikzad, N.; Chenaghlu, M.; Socher, R.; Amatriain, X.; and Gao, J. 2024. Large language models: A survey.
- Mitra, A., and Baral, C. 2016. Addressing a question answering challenge by combining statistical methods with in-

ductive rule learning and reasoning. In *AAAI*, 2779–2785. AAAI Press.

Naveed, H.; Khan, A. U.; Qiu, S.; Saqib, M.; Anwar, S.; Usman, M.; Akhtar, N.; Barnes, N.; and Mian, A. 2024. A comprehensive overview of large language models.

Nye, M. I.; Tessler, M. H.; Tenenbaum, J. B.; and Lake, B. M. 2021. Improving coherence and consistency in neural sequence models with dual-system, neuro-symbolic reasoning. In *NeurIPS*, 25192–25204.

Peng, S.; Kalliamvakou, E.; Cihon, P.; and Demirer, M. 2023. The impact of AI on developer productivity: Evidence from github copilot. *CoRR* abs/2302.06590.

Raiaan, M. A. K.; Mukta, M. S. H.; Fatema, K.; Fahad, N. M.; Sakib, S.; Mim, M. M. J.; Ahmad, J.; Ali, M. E.; and Azam, S. 2024. A review on large language models: Architectures, applications, taxonomies, open issues and challenges. *IEEE Access* 12:26839–26874.

Rajasekharan, A.; Zeng, Y.; Padalkar, P.; and Gupta, G. 2023. Reliable natural language understanding with large language models and answer set programming. *Electronic Proceedings in Theoretical Computer Science* 385:274–287.

Schwitter, R. 2018. Specifying and verbalising answer set programs in controlled natural language. *Theory Pract. Log. Program.* 18(3-4):691–705.

Stahlberg, F. 2020. Neural machine translation: A review. *J. Artif. Intell. Res.* 69:343–418.

Team, G. 2024a. Gemini: A family of highly capable multimodal models.

Team, G. 2024b. Gemma: Open models based on gemini research and technology.

Touvron, H., et al. 2023. Llama 2: Open foundation and fine-tuned chat models.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is all you need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

Wang, Y.; Le, H.; Gotmare, A. D.; Bui, N. D. Q.; Li, J.; and Hoi, S. C. H. 2023. Codet5+: Open code large language models for code understanding and generation.

Xu, F. F.; Alon, U.; Neubig, G.; and Hellendoorn, V. J. 2022. A systematic evaluation of large language models of code. In *Proceedings of the 6th ACM SIGPLAN International Symposium on Machine Programming*, MAPS 2022, 1–10. New York, NY, USA: Association for Computing Machinery.

Yang, Z.; Ishay, A.; and Lee, J. 2023. Coupling large language models with logic programming for robust and general reasoning from text. In *ACL (Findings)*, 5186–5219. ACL.

Ziegler, D. M.; Stiennon, N.; Wu, J.; Brown, T. B.; Radford, A.; Amodei, D.; Christiano, P.; and Irving, G. 2020. Fine-tuning language models from human preferences.