

Kara: A System for Visualising and Visual Editing of Interpretations for Answer-Set Programs

Christian Kloimüller¹, Johannes Oetsch², Jörg Pührer^{2(✉)}, and Hans Tompits²

¹ Forschungsgruppe für Industrielle Software (INSO),
Technische Universität Wien,

Favoritenstraße 9-11, 1040 Vienna, Austria
`christian.kloimuellner@inso.tuwien.ac.at`

² Institut für Informationssysteme 184/3,
Technische Universität Wien,
Favoritenstraße 9-11, 1040 Vienna, Austria
`{oetsch, puehrer, tompits}@kr.tuwien.ac.at`

Abstract. In answer-set programming (ASP), the solutions of a problem are encoded in dedicated models, called *answer sets*, of a logical theory. These answer sets are computed from the program that represents the theory by means of an ASP solver and returned to the user as sets of ground first-order literals. As this type of representation is often cumbersome for the user to interpret, tools like `ASPVIZ` and `IDPDraw` were developed that allow for visualising answer sets. The tool `Kara`, introduced in this paper, follows these approaches, using ASP itself as a language for defining visualisations of interpretations. Unlike existing tools that position graphic primitives according to static coordinates only, `Kara` allows for more high-level specifications, supporting graph structures, grids, and relative positioning of graphical elements. Moreover, generalising the functionality of previous tools, `Kara` provides modifiable visualisations such that interpretations can be manipulated by graphically editing their visualisations. This is realised by resorting to abductive reasoning techniques using ASP itself. `Kara` is part of `SeaLion`, an integrated development environment (IDE) for ASP.

1 Introduction

Answer-set programming (ASP) [1] is a well-known paradigm for declarative problem solving. Its key idea is that a problem is encoded in terms of a logic program such that dedicated models of it, called *answer sets*, correspond to the solutions of the problem. Answer sets are interpretations, usually represented by sets of ground first-order literals.

A problem often faced when developing answer-set programs is that interpretations returned by an ASP solver are cumbersome to read—in particular, in

This work was partially supported by the Austrian Science Fund (FWF) under project P21698.

case of large interpretations which are spread over several lines on the screen or the output file. Hence, a user may have difficulties extracting the relevant information from the textual representation of an answer set. Related to this issue, there is one even harder practical problem: editing or writing interpretations by hand.

Although the general goal of ASP is to have answer sets computed automatically, we identify different situations during the development of answer-set programs in which it would be helpful to have adequate means to manipulate interpretations. First, in declarative debugging [2], the user has to specify the expected semantics in order for the debugging system to identify the causes for a mismatch with the actual semantics. In previous work [3], a debugging approach has been introduced that takes a program P and an interpretation I that is expected to be an answer set of P and returns reasons why I is not an answer set of P . Manually producing such an intended interpretation ahead of computation is a time-consuming task, however. Another situation in which the creation of an interpretation can be useful is testing post-processing tools. Typically, if answer-set solvers are used within other applications, they are embedded as a module in a larger context. The overall application delegates a problem to the solver by transforming it to a respective answer-set program, and the outcome of the solver is then processed further as needed by the application. In order to test post-processing components, which may be written by programmers unaware of ASP, it would be beneficial to have means to create mock answer sets as test inputs. Third, the same idea of providing test input applies to modular answer-set programming [4], when a module B that depends on another module A is developed before or separately from A . To test B , B can be joined with interpretations mocking answer sets from A .

In this paper, we describe the system **Kara** which allows for both visualising interpretations and editing them by manipulating their visualisations.¹ The visualisation functionality of **Kara** has been inspired by the existing tools **ASPVIZ** [5] and **IDPDraw** [6] for visualising answer sets. The key idea is to use ASP itself as a language for specifying how to visualise an interpretation. To this end, the user takes a dedicated answer-set program V —which we call a *visualisation program*—that specifies how the visualisation of an interpretation I should look like. That is, V defines how different graphical elements, such as rectangles, polygons, images, graphs, etc., should be arranged and configured to visually represent I .

Kara offers a rich visualisation language that allows for defining a super-set of the graphical elements available in **ASPVIZ** and **IDPDraw**, e.g., providing support for automatically layouting graph structures, relative and absolute positioning, and support for grids of graphical elements. Moreover, **Kara** also offers a *generic mode* of visualisation, not available in previous tools, that does not require a domain-specific visualisation program, and visualises an answer set as

¹ The name “**Kara**” derives, with all due respect, from “Kara Zor-El”, the native Kryptonian name of *Supergirl*, given that Kryptonians have visual superpowers on Earth.




Fig. 1. Overview of the workflow (visualisation and abduction process).

a hypergraph whose set of nodes corresponds to the individuals occurring in the interpretation. A detailed overview of the differences concerning the visualisation capabilities of *Kara* with other tools is given in Sect. 5. A general difference to previous tools is that *Kara* does not just produce image files right away but presents the visualisation in form of modifiable graphical elements in a visual editor. The user can manipulate the visualisation in various ways, e.g., change size, position, or other properties of graphical elements, as well as copy, delete, and insert new ones. Notably, the created visualisations can also be used outside our editing framework as *Kara* offers an SVG export function that allows to save the possibly modified visualisation as a vector graphic. Besides fine-tuning exported SVG files, manipulation of the visualisation of an interpretation I can be done for obtaining a modified version I' of I by means of ASP-based abductive reasoning [7]. This gives the possibility to visually edit interpretations which is useful for debugging and testing purposes as described above. We will present a number of examples to illustrate the functionality of *Kara* and the ease of coping with a visualised answer set compared to interpreting its textual representation.

Kara is designed as a plugin of *SeaLion*, an Eclipse-based integrated development environment (IDE) for ASP [8] that is currently developed as part of a project on programming-support methods for ASP [9].

2 System Overview

We assume familiarity with the basic concepts of ASP (for a thorough introduction to the subject, cf. Baral [1]). In brief, an answer-set program consists of rules of the form

$$a_1 \vee \dots \vee a_l : - a_{l+1}, \dots, a_m, \text{not } a_{m+1}, \dots, \text{not } a_n,$$

where $n \geq m \geq l \geq 0$, “not” denotes *default negation*, and all a_i are first-order literals (i.e., atoms possibly preceded by the *strong negation* symbol \neg). For a rule r as above, we define the *head* of r as $H(r) = \{a_1, \dots, a_l\}$ and the *body* as $B(r) = \{a_{l+1}, \dots, a_m, \text{not } a_{m+1}, \dots, \text{not } a_n\}$. We will also refer to the *positive body*, given as $B^+(r) = \{a_{l+1}, \dots, a_m\}$. If $n = l = 1$, r is a *fact*, and if $l = 0$, r is a

constraint. For facts, we will usually omit the symbol “:–”. The *grounding* of a program P relative to its Herbrand universe is defined as usual. An *interpretation* I is a finite and consistent set of ground literals, where consistency means that $\{a, \neg a\} \not\subseteq I$, for any atom a . An interpretation I is an *answer set* of a program P if it is a subset-minimal model of the grounding of the *reduct* of P relative to I (see Baral [1] for details).

The overall workflow of *Kara* is depicted in Fig. 1, illustrating how an interpretation I can be visualised in the upper row and how changing the visualisation can be reflected back to I such that we obtain a modified version I' of I in the lower row. In the following, we call programs that encode problems for which I and I' represent solution candidates *domain programs*.

2.1 Visualisation of Interpretations

As discussed in the introduction, we use ASP itself as a language for specifying how to visualise an interpretation. In doing so, we follow a similar approach as the tools ASPVIZ [5] and IDPDraw [6]. We next describe this method on an abstract level.

Assume we want to visualise an interpretation I that is defined over a first-order alphabet \mathcal{A} . We join I , interpreted as a set of facts, with a visualisation program V that is defined over $\mathcal{A}' \supset \mathcal{A}$, where \mathcal{A}' may contain auxiliary predicates and function symbols, as well as predicates from a fixed set \mathcal{P}_v of reserved *visualisation predicates* that vary for the three tools.²

The rules in V are used to derive different atoms with predicates from \mathcal{P}_v , depending on I , that control the individual graphical elements of the resulting visualisation including their presence or absence, position, and all other properties. An actual visualisation is obtained by post-processing an answer set I_v of $V \cup I$ that is projected to the predicates in \mathcal{P}_v . We refer to I_v as a *visualisation answer set* for I . Note that since V is an arbitrary answer-set program it might be non-deterministic in the sense that multiple visualisation answer sets may exist. In the current implementation only one of them is used for visualisation. The process is depicted in the upper row of Fig. 1. An exhaustive list of visualisation predicates available in *Kara* is given in Appendix A.

Example 1. Assume we deal with a domain program whose answer sets correspond to arrangements of items on two shelves. Consider the interpretation

$$I = \{book(s_1, 1), book(s_1, 3), book(s_2, 1), globe(s_2, 2)\}$$

stating that two books are located on shelf s_1 in positions 1 and 3 and that there is another book and a globe on shelf s_2 in positions 1 and 2, respectively. The goal is to create a simple graphical representation of this and similar interpretations,

² Technically, in ASPVIZ, V is not joined with I but with a domain program P such that I is an answer set of P .




Fig. 2. The visualisation of interpretation I from Example 1.

depicting the two shelves as two lines, each book as a rectangle, and globes as circles. Consider the following visualisation program:

$$\text{visline}(\text{shelf}_1, 10, 40, 80, 40, 0), \quad (1)$$

$$\text{visline}(\text{shelf}_2, 10, 80, 80, 80, 0), \quad (2)$$

$$\text{visrect}(f(X, Y), 20, 8) : - \text{book}(X, Y), \quad (3)$$

$$\text{visposition}(f(s_1, Y), 20 * Y, 20, 0) : - \text{book}(s_1, Y), \quad (4)$$

$$\text{visposition}(f(s_2, Y), 20 * Y, 60, 0) : - \text{book}(s_2, Y), \quad (5)$$

$$\text{visellipse}(f(X, Y), 20, 20) : - \text{globe}(X, Y), \quad (6)$$

$$\text{visposition}(f(s_1, Y), 20 * Y, 20, 0) : - \text{globe}(s_1, Y), \quad (7)$$

$$\text{visposition}(f(s_2, Y), 20 * Y, 60, 0) : - \text{globe}(s_2, Y). \quad (8)$$

Rules (1) and (2) create two lines with the identifiers shelf_1 and shelf_2 , representing the top and bottom shelf. The second to fifth arguments of $\text{visline}/6$ represent the origin and the target coordinates of the line.³ The last argument of $\text{visline}/6$ is a z -coordinate determining which graphical element is visible in case two or more overlap. Rule (3) generates the rectangles representing books, and Rules (4) and (5) determine their position depending on the shelf and the position given in the interpretation. Likewise, Rules (6) to (8) generate and position globes. The resulting visualisation of I is depicted in Fig. 2. \square

Note that the first argument of each visualisation predicate is a unique identifier for the respective graphical element. By making use of function symbols with variables, like $f(X, Y)$ in Rule (3) above, these labels are not limited to constants in the visualisation program but can be generated on the fly, depending on the interpretation to visualise. While some visualisation predicates, like visline , visrect , and visellipse , define graphical elements, others, e.g., visposition , are used to change properties of the elements, referring to them by their respective identifiers.

Kara also offers a *generic visualisation* that visualises an arbitrary interpretation without the need for defining a visualisation program. In such a case, the interpretation is represented as a labelled hypergraph. Its nodes are the

³ The origin of the coordinate system is at the top-left corner of the illustration window with the x -axis pointing to the right and the y -axis pointing down.

individuals appearing in the interpretation and the edges represent the literals in the interpretation, connecting the individuals appearing in the respective literal. Integer labels on the endings of the edge are used for expressing the term position of the individual. To distinguish between different predicates, each edge has an additional label stating the predicate. Edges of the same predicate are of the same colour. A generic visualisation is presented in Example 4 in Sect. 4.

2.2 Editing of Interpretations

We next describe how we can obtain a modified version I' of an interpretation I corresponding to a manipulation of the visualisation of I . We follow the steps depicted in the lower row of Fig. 1, using abductive reasoning. Recall that abduction is the process of finding hypotheses that explain given observations in the context of a theory. Intuitively, in our case, the theory is the visualisation program, the observation is the modified visualisation of I , and the desired hypothesis is I' .

In *Kara*, the visualisation of I is created using the Graphical Editing Framework (GEF) [10] of Eclipse. It is displayed in a graphical editor which allows for various kinds of manipulation actions such as moving, resizing, adding or deleting graphical elements, adding or removing edges between them, editing their properties, or changing grid values. Each change in the visual editor of *Kara* is internally reflected by a modification to the underlying visualisation answer set I_v . We denote the visualisation interpretation that results from editing I_v as I'_v . From that and the visualisation program V , we construct a logic program $\lambda(I'_v, V)$ such that the visualisation of any answer set I' of $\lambda(I'_v, V)$ using V corresponds to the modified one.

The idea is that $\lambda(I'_v, V)$, which we refer to as the *abduction program* for I'_v and V , guesses a set of *abducible atoms*. On top of these atoms, the rules of V are used in $\lambda(I'_v, V)$ to derive a hypothetical visualisation answer set I''_v for I' . Finally, constraints in the abduction program ensure that I''_v coincides with the targeted visualisation interpretation I'_v on a set \mathcal{P}_i of selected predicates from \mathcal{P}_v , which we call *integrity predicates*. Hence, a modified interpretation I' can be obtained by computing an answer set of $\lambda(I'_v, V)$ and projecting it to the guessed atoms. To summarise, the abduction problem underlying the described process can be stated as follows:

- (*) Given the interpretation I'_v , determine an interpretation I' such that I'_v coincides with each answer set of $V \cup I'$ on \mathcal{P}_i .

Naturally, depending on V and I'_v it is possible that no such solution I' exists. Visualisation programs must be written in a way such that manipulated visualisation interpretations could indeed be the outcome of the visualisation program for some input. This is not the case for arbitrary visualisation programs, but usually it is easy to write an appropriate visualisation program that allows for abducting interpretations.

$$\begin{aligned}
\text{dom}(I'_v, V) = & \{ \text{nonRecDom}(t) : - v(\mathbf{t}') \mid r \in V, v/m \in \mathcal{P}_v, v(\mathbf{t}') \in H(r), \\
& a(\mathbf{t}) \in B^+(r), \mathbf{t} = t_1, \dots, t, \dots, t_n, a/n \notin \mathcal{P}_v, \\
& \text{VAR}(t) \neq \emptyset, \text{VAR}(t) \subseteq \text{VAR}(\mathbf{t}') \} \cup \\
& \{ \text{dom}(t) : - v(\mathbf{t}'), \text{nonRecDom}(X_1), \dots, \text{nonRecDom}(X_l) \mid r \in V, \\
& v/m \in \mathcal{P}_v, v(\mathbf{t}') \in H(r), a(\mathbf{t}) \in B^+(r), \mathbf{t} = t_1, \dots, t, \dots, t_n, \\
& a/n \notin \mathcal{P}_v, \text{VAR}(t) \cap \text{VAR}(\mathbf{t}') \neq \emptyset, \\
& \text{VAR}(t) \setminus \text{VAR}(\mathbf{t}') = \{X_1, \dots, X_l\} \} \cup \\
& \{ \text{dom}(X) : - \text{nonRecDom}(X) \}. \\
\text{guess}(V) = & \{ a(X_1, \dots, X_n) : - \text{not } \neg a(X_1, \dots, X_n), \text{dom}(X_1), \dots, \text{dom}(X_n), \\
& \neg a(X_1, \dots, X_n) : - \text{not } a(X_1, \dots, X_n), \text{dom}(X_1), \dots, \text{dom}(X_n) \mid \\
& a/n \notin \mathcal{P}_v, a(t_1, \dots, t_n) \in \bigcup_{r \in V} B(r), \\
& \{ a(t'_1, \dots, t'_n) \mid a(t'_1, \dots, t'_n) \in H(r), r \in V \} = \emptyset \}. \\
\text{check}(I'_v) = & \{ : - \text{not } v(t_1, \dots, t_n), : - v(X_1, \dots, X_n), \text{not } v'(X_1, \dots, X_n), \\
& v'(t_1, \dots, t_n) \mid v(t_1, \dots, t_n) \in I'_v, v/n \in \mathcal{P}_i \}.
\end{aligned}$$

Fig. 3. Elements of the abduction program $\lambda(I'_v, V)$.

The following problems have to be addressed for realising the sketched approach:

- determining the predicates and domains of the abducible atoms, and
- choosing the integrity predicates among the visualisation predicates.

For solving these issues, we rely on pragmatic choices that seem useful in practice. We obtain the set \mathcal{P}_a of predicates of the abducible atoms from the visualisation program V . The idea is that every predicate that is relevant to the solution of a problem encoded in an answer set has to occur in the visualisation program if the latter is meant to provide a complete graphical representation of the solution. Moreover, we restrict \mathcal{P}_a to those non-visualisation predicates in V that occur in the body of a rule but not in any head atom in V . The assumption is that atoms defined in V are most likely of auxiliary nature and not contained in a domain program.

An easy approach for generating a domain \mathcal{D}_a of the abducible atoms would be to extract the terms occurring in I'_v . We follow, however, a more fine-grained approach that takes the introduction and deletion of function symbols in the rules in V into account. Assume V contains the rules

$$\begin{aligned}
& \text{visrect}(f(\text{Street}, \text{Num}), 9, 10) : - \text{house}(\text{Street}, \text{Num}) \quad \text{and} \\
& \text{visellipse}(\text{sun}, \text{Width}, \text{Height}) : - \text{property}(\text{sun}, \text{size}(\text{Width}, \text{Height})),
\end{aligned}$$

and I'_v contains $\text{visrect}(f(\text{bakerstreet}, 221b), 9, 10)$ and $\text{visellipse}(\text{sun}, 10, 11)$. Then, when extracting the terms in I'_v , the domain includes $f(\text{bakerstreet}, 221b)$, bakerstreet , $221b$, 9 , 10 , sun , and 11 for the two rules. However, the functor f is solely an auxiliary concept in V and not meant to be part of domain programs.

Moreover, the term 9 is introduced in V and is not needed in the domain for I' . Also, the terms 10 and 11 as standalone terms and *sun* are not needed in I' to derive I'_v . Even worse, the term $size(10, 11)$, that has to be contained in I' such that I'_v can be a visualisation answer set for I' , is missing in the domain. Hence, we derive \mathcal{D}_a in $\lambda(I'_v, V)$ not only from I'_v but also consider the rules in V . Using our translation detailed below, we obtain *bakerstreet*, *221b*, and $size(10, 12)$ as domain terms from the rules above.

For the choice of \mathcal{P}_i , i.e., of the predicates on which I'_v and the actual visualisation answer sets of I' need to coincide, we exclude visualisation predicates that require a high preciseness in visual editing by the user in order to match exactly a value that could result from the visualisation program. For example, we do not include predicates determining position and size of graphical elements, since in general it is hard to position and scale an element precisely such that an interpretation I' exists with a matching visualisation. Note that this is not a major restriction, as in general it is easy to write a visualisation program such that aspects that the user wants to be modifiable are represented by graphical elements that can be elegantly modified visually. For example, instead of representing a Sudoku puzzle by labels whose exact position is calculated in the visualisation program, the language of *Kara* allows for using a logical grid such that the value of each cell can be easily changed in the visual editor.

We next give the details of the abduction program.

Definition 1. Let I'_v be an interpretation over predicates in \mathcal{P}_v , V a (visualisation) program, and $\mathcal{P}_i \subseteq \mathcal{P}_v$ the fixed set of integrity predicates. Moreover, let $VAR(T)$ denote the variables occurring in T , where T is a term or a list of terms. Then, the abduction program $\lambda(I'_v, V)$ with respect to I'_v and V is given by

$$\lambda(I'_v, V) = \text{dom}(I'_v, V) \cup \text{guess}(V) \cup V \cup \text{check}(I'_v),$$

where $\text{dom}(I'_v, V)$, $\text{guess}(V)$, and $\text{check}(I'_v)$ are given in Fig. 3, and $\text{nonRecDom}/1$, $\text{dom}/1$, and v'/n , for all $v/n \in \mathcal{P}_i$, are fresh predicates.

The idea of $\text{dom}(I'_v, V)$ is to consider non-ground terms t contained in the body of a visualisation rule that share variables with a visualisation atom in the head of the rule and to derive instances of these terms when the corresponding visualisation atom is contained in I'_v . In case less variables occur in the visualisation atom than in t , we avoid safety problems by restricting their scope to parts of the derived domain. Here, the distinction between predicates dom and nonRecDom is necessary to prevent infinite groundings of the abduction program. The next part of the abduction program is $\text{guess}(V)$, where the atoms of the domain program P are guessed, i.e., the abducible atoms. The output of the guessing part is I' . Finally, $\text{check}(I'_v)$ contains all constraints and auxiliary facts.

Note that in general it is not guaranteed that the domain we derive contains all necessary elements for abducting an appropriate interpretation I' . For instance, consider the case that the visualisation program contains a rule

$$\text{visrect}(id, 5, 5) :- \text{foo}(X),$$

and V , together with the constraints in $\text{check}(I'_v)$, require that for every term t of a domain that can be obtained from I'_v and V , $\text{foo}(t)$ must not hold. Then, there is no interpretation that will trigger the rule using this domain, although an interpretation with a further term t' might exist that results in the desired visualisation. Hence, we added an editor to Kara that allows for changing and extending the automatically generated domain as well as the set of abducible predicates.

The following two results characterise the answer sets of the abduction program.

Theorem 1. *Let I'_v be an interpretation with atoms over predicates in \mathcal{P}_v , V a (visualisation) program, and $\mathcal{P}_i \subseteq \mathcal{P}_v$ the fixed set of integrity predicates. Then, any answer set I''_v of $\lambda(I'_v, V)$ coincides with I'_v on the atoms over predicates from \mathcal{P}_i .*

Note that since the stated abduction problem (*) requires I'_v to coincides with each answer set of $V \cup I'$ on \mathcal{P}_i , a solution is only given in case the visualisation program deterministically derives a visualisation, as expressed in Theorem 2.

Theorem 2. *Let I'_v be an interpretation with atoms over \mathcal{P}_v , $\mathcal{P}_i \subseteq \mathcal{P}_v$ the fixed set of integrity predicates, and V a (visualisation) program such that, for every I' with atoms over \mathcal{P}_d , where*

$$\mathcal{P}_d = \{a/n \mid \text{there are terms } t_1, \dots, t_n \text{ such that } a(t_1, \dots, t_n) \in \bigcup_{r \in V} B(r) \text{ but there are no terms } t'_1, \dots, t'_n \text{ such that } a(t'_1, \dots, t'_n) \in \bigcup_{r \in V} H(r) \} \setminus \mathcal{P}_v,$$

every two answer sets I_1 and I_2 of $V \cup I'$ do not differ on \mathcal{P}_i . Then, for any answer set I''_v of $\lambda(I'_v, V)$, a solution I' of the abduction problem () is obtained by projecting I''_v to the predicates in \mathcal{P}_d .*

3 Implementation

Kara is written in Java and integrated in the Eclipse-plugin SeaLion [8] for developing answer-set programs. SeaLion is an IDE that offers functionality




Fig. 4. Technology stack of the Kara system.

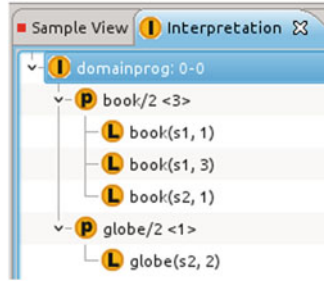


Fig. 5. Sample output in the *interpretation view* of SeaLion.

to execute external ASP solvers on answer-set programs and thus realises the interface between Kara and ASP solvers. The overall technology stack of Kara is depicted in Fig. 4. Currently, programs in the languages of Gringo and DLV are supported.

Next we describe how to use Kara in SeaLion. The ASP developer may invoke ASP solver calls in SeaLion using Eclipse’s launch-configuration framework in a similar fashion as Java programs are started from within Eclipse. Answer sets resulting from a solver call can be parsed by the IDE and displayed as expandable tree structures in a dedicated Eclipse view for interpretations as shown in Fig. 5. Starting from there, the user can invoke Kara by choosing a pop-up menu entry of the interpretation that should be visualised. Here, one can select between the generic mode of visualisation or a mode using a visualisation program created by the user. In the latter case, a customised run configuration dialog will open that allows for choosing a file containing the visualisation program and for setting the solver configuration, including selection of the solver and command-line arguments, to be used by Kara. Then, the visual editor opens with the generated visualisation. It allows for changing the visualisation in various ways, including repositioning, rescaling, and renaming of graphical elements, as well as creation, deletion, and making copies of elements. The current state of the visualisation can always be exported in an SVG file. Moreover, the process for abducting an interpretation that reflects the modifications to the visualisation can be started from a pop-up menu of the visual editor. If a respective interpretation exists, it will be added to the interpretation view of SeaLion.

Kara and SeaLion can be downloaded and installed from within Eclipse, using the following update site:

<http://sealion.at/update>.

For more information and installation instructions, we refer to the project web site

<http://www.sealion.at>.

```

visgrid(maze, MR, MC, MR*20+5, MC*20+5):- maxC(MC), maxR(MR). (9)
visposition(maze, 0, 0, 0). (10)
% A cell with a wall on it.
visrect(wall, 20, 20). (11)
visbackgroundcolor(wall, black). (12)
% An empty cell.
visrect(empty, 20, 20). (13)
visbackgroundcolor(empty, white). (14)
viscolor(empty, white). (15)
% Entrance and exit.
visimage(entrance, "entrance.jpg"). (16)
visscale(entrance, 18, 18). (17)
visimage(exit, "exit.png"). (18)
visscale(exit, 18, 18). (19)
% Filling the cells of the grid.
visfillgrid(maze, empty, R, C):- empty(C, R), not entrance(C, R),
                                not exit(C, R). (20)
visfillgrid(maze, wall, R, C):- wall(C, R), not entrance(C, R),
                                not exit(C, R). (21)
visfillgrid(maze, entrance, R, C):- entrance(C, R). (22)
visfillgrid(maze, exit, R, C):- exit(C, R). (23)
% Vertical and horizontal lines.
visline(v(0), 5, 5, 5, MR * 20 + 5, 1):- maxR(MR). (24)
visline(v(C), C*20+5, 5, C*20+5, MR*20+5, 1):- col(C), maxR(MR). (25)
visline(h(0), 5, 5, MC * 20 + 5, 5, 1):- maxC(MC). (26)
visline(h(R), 5, R * 20 + 5, MC * 20 + 5, R * 20 + 5, 1):- row(R),
                                                            maxC(MC). (27)
% Define possible grid values for editing.
vispossiblegridvalues(maze, wall). (28)
vispossiblegridvalues(maze, empty). (29)
vispossiblegridvalues(maze, entrance). (30)
vispossiblegridvalues(maze, exit). (31)

```

Fig. 6. Visualisation program for Example 2.

4 Examples

In this section, we provide examples that give an overview of the functionality of Kara. We first illustrate the use of logic grids and the visual editing feature.

Example 2. *Maze generation* is a benchmark problem from the second ASP competition [11]. The task is to generate a two-dimensional grid, where each cell is

either a wall or empty, that satisfies certain constraints. There are two dedicated empty cells, being the maze's entrance and its exit, respectively. The following facts represent a sample answer set of a maze-generation encoding restricted to interesting predicates:

```
col(1..5). row(1..5). maxC(5). maxR(5). wall(1,1). empty(1,2). wall(1,3).
wall(1,4). wall(1,5). wall(2,1). empty(2,2). empty(2,3). empty(2,4).wall(2,5).
wall(3,1). wall(3,2). wall(3,3). empty(3,4). wall(3,5). wall(4,1). empty(4,2).
empty(4,3). empty(4,4). wall(4,5). wall(5,1).wall(5,2). wall(5,3). empty(5,4).
wall(5,5). entrance(1,2). exit(5,4).
```

Predicates *col/1* and *row/1* define indices for the rows and columns of the maze, while *maxC/1* and *maxR/1* give the maximum column and row index, respectively. The predicates *wall/2*, *empty/2*, *entrance/2*, and *exit/2* determine the positions of walls, empty cells, the entrance, and the exit in the grid, respectively. One may use the visualisation program from Fig. 6 for maze-generation interpretations of this kind.

In Fig. 6, Rule (9) defines a logic grid with identifier *maze* having *MR* rows and *MC* columns. The fourth and fifth parameter define the height and width of the grid in pixel. Rule (10) is a fact that defines a fixed position for the maze. The next step is to define the graphical objects to be displayed in the grid. Because these objects are fixed (i.e., they are used more than once), they can be defined as facts. A wall is represented by a rectangle with black background and foreground colour⁴ (Rules (11) and (12)) whereas an empty cell is rendered as a rectangle with white background and foreground colour (Rules (13) to (15)). The entrance and the exit are represented by two images (Rules (16) to (19)). Then, these graphical elements are assigned to the respective cell of the grid (Rules (20) to (23)). Rules (24) to (27) render vertical and horizontal lines to better distinguish between the different cells. Rules (28) to (31) are not needed for visualisation but define possible values for the grid that we want to be available in the visual editor.

Once the grid is rendered, the user can replace the value of a cell with a value defined by using predicate *vispossiblegridvalues/2* (e.g., replacing an empty cell




Fig. 7. Visualisation output for the maze-generation program.

⁴ Black foreground colour is default and need not be set explicitly.

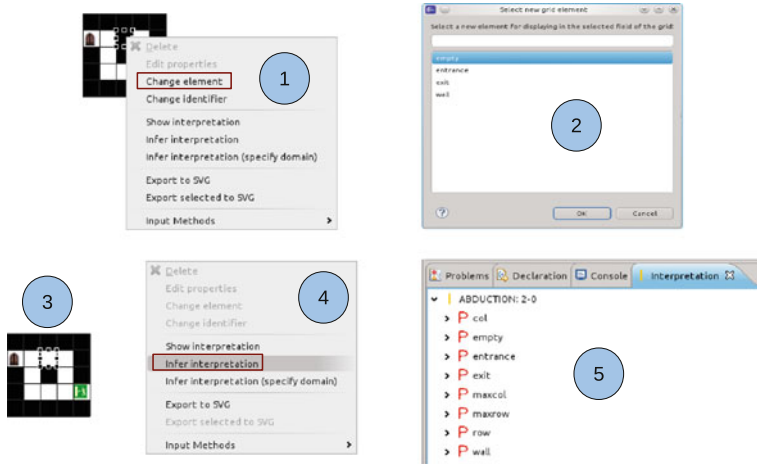


Fig. 8. Abduction steps in the plugin.

with a wall). The visualisation of the sample interpretation using this program is depicted in Fig. 7. Evidently, the visual representation of the answer set is more accessible than the textual representation of the answer set given in the beginning of the example.

Next, we demonstrate how to use the visual editing feature of Kara to obtain a modified interpretation; the respective steps are illustrated by Fig. 8. Suppose we want to change the cell (3, 2) from being a wall to an empty cell. The user can select the respective cell and open a pop-up menu that provides an item for changing grid-values. A dialog opens that allows for choosing among the values that have been defined in the visualisation program, using the *vispossiblegridvalues/2* predicate. When the user has finished editing the visualisation, the abduction process for inferring the new interpretation can be started. After an interpretation is derived, it is added to SeaLion's interpretation view. \square

Kara supports absolute and relative positioning of graphical elements. If for any visualisation element the predicate *visposition/4* is defined, then fixed positioning is used. Otherwise, the element is positioned automatically. Then, by default, the elements are randomly positioned on the graphical editor. However, the user can define the position of an element *relative* to another element. This is done by using the predicates *visleft/2*, *visright/2*, *visabove/2*, *visbelow/2*, and *visinfrontof/2*.

Example 3. The following visualisation program makes use of relative positioning for sorting elements according to their label.

$$\text{visrect}(a, 50, 50). \quad (32)$$

$$\text{vislabel}(a, \text{label}). \quad (33)$$

$$\text{vistext}(\text{label}, 3). \quad (34)$$

$$\text{vispolygon}(b, 0, 20, 1). \quad (35)$$

$$\text{vispolygon}(b, 25, 0, 2). \quad (36)$$

$$\text{vispolygon}(b, 50, 20, 3). \quad (37)$$

$$\text{vislabel}(b, \text{labb}). \quad (38)$$

$$\text{vistext}(\text{labb}, 10). \quad (39)$$

$$\text{visellipse}(c, 30, 30). \quad (40)$$

$$\text{vislabel}(c, \text{labc}). \quad (41)$$

$$\text{vistext}(\text{labc}, 5). \quad (42)$$

$$\text{element}(X) :- \text{visrect}(X, -, -). \quad (43)$$

$$\text{element}(X) :- \text{vispolygon}(X, -, -, -). \quad (44)$$

$$\text{element}(X) :- \text{visellipse}(X, -, -). \quad (45)$$

$$\text{visleft}(X, Y) :- \text{element}(X), \text{element}(Y), \text{vislabel}(X, \text{LABX}), \quad (46)$$

$$\text{vistext}(\text{LABX}, \text{XNUM}), \text{vislabel}(Y, \text{LABY}),$$

$$\text{vistext}(\text{LABY}, \text{YNUM}), \text{XNUM} < \text{YNUM}.$$

The program defines three graphical objects, a rectangle, a polygon, and an ellipse. In Rules (32) to (34), the rectangle together with its label, 3, is generated. The shape of the polygon (Rules (35) to (37)) is defined by a sequence of points relative to the polygon's own coordinate system using the *vispolygon*/4 predicate. The order in which these points are connected with each other is given by the predicate's fourth argument. Rules (38) and (39) generate the label for the polygon and specify its text. Rules (43) to (45) state that every rectangle, polygon, and ellipse is an element. The relative position of the three elements is determined by Rule (46). For two elements E_1 and E_2 , E_1 has to appear to the left of E_2 whenever the label of E_1 is smaller than the one of E_2 .

The output of this visualisation program is shown in Fig. 9. Note that the visualisation program does not make reference to predicates from an interpretation to visualise, hence the example illustrates that *Kara* can also be used for creating arbitrary graphics. \square

The last example demonstrates the support for graphs in *Kara*. Moreover, the generic visualisation feature is illustrated.

Example 4. Suppose we want to visualise answer sets of an encoding of a graph-colouring problem. Assume we have the following interpretation that defines nodes and edges of a graph as well as a colour for each node:

$$\{\text{node}(1), \text{node}(2), \text{node}(3), \text{node}(4), \text{node}(5), \text{node}(6), \text{edge}(1, 2), \text{edge}(1, 3), \\ \text{edge}(1, 4), \text{edge}(2, 4), \text{edge}(2, 5), \text{edge}(2, 6), \text{edge}(3, 1), \text{edge}(3, 4), \text{edge}(3, 5), \\ \text{edge}(4, 1), \text{edge}(4, 2), \text{edge}(5, 3), \text{edge}(5, 4), \text{edge}(5, 6), \text{edge}(6, 2), \text{edge}(6, 3), \\ \text{edge}(6, 5), \text{colour}(1, \text{lightblue}), \text{colour}(2, \text{yellow}), \text{colour}(3, \text{yellow}), \\ \text{colour}(4, \text{red}), \text{colour}(5, \text{lightblue}), \text{colour}(6, \text{red})\}.$$




Fig. 9. Output of the visualisation program in Example 3.

We make use of the following visualisation program:

```
% Generate a graph.
visgraph(g). (47)
```

```
% Generate the nodes of the graph.
visellipse(X, 20, 20) :- node(X). (48)
```

```
visinode(X, g) :- node(X). (49)
```

```
% Connect the nodes (edges of the input).
visconnect(f(X, Y), X, Y) :- edge(X, Y). (50)
```

```
visargetdeco(X, arrow) :- visconnect(X, -, -). (51)
```

```
% Generate labels for the nodes.
vislabel(X, l(X)) :- node(X). (52)
```

```
visetext(l(X), X) :- node(X). (53)
```

```
visfontstyle(l(X), bold) :- node(X). (54)
```

```
% Colour the node according to the solution.
visbackgroundcolor(X, COL) :- node(X), color(X, COL). (55)
```

In Rule (47), a graph, g , is defined and a circle for every node from the input interpretation is created (Rule (48)). Rule (49) states that each of these circles is logically considered a node of graph g . This has the effect that they will be considered by the algorithm layouting the graph during the creation of the visualisation. The edges of the graph are defined using the *visconnect/3* predicate (Rule (50)). It can be used to connect arbitrary graphical elements with a line, also if they are not nodes of some graph. As we deal with a directed graph, an arrow is set as target decoration for all the connections (Rule (51)). Labels for the nodes are set in Rules (52) to (54). Finally, Rule (55) sets the colour of the node according to the interpretation. The resulting visualisation is depicted in Fig. 10. Moreover, the generic visualisation of the graph colouring interpretation is given in Fig. 11. \square

5 Related Work

The visualisation method realised in *Kara* follows the approach of the previous systems *ASPVIZ* [5] and *IDPDraw* [6], which also use ASP for defining how




Fig. 10. Visualisation of a coloured graph.

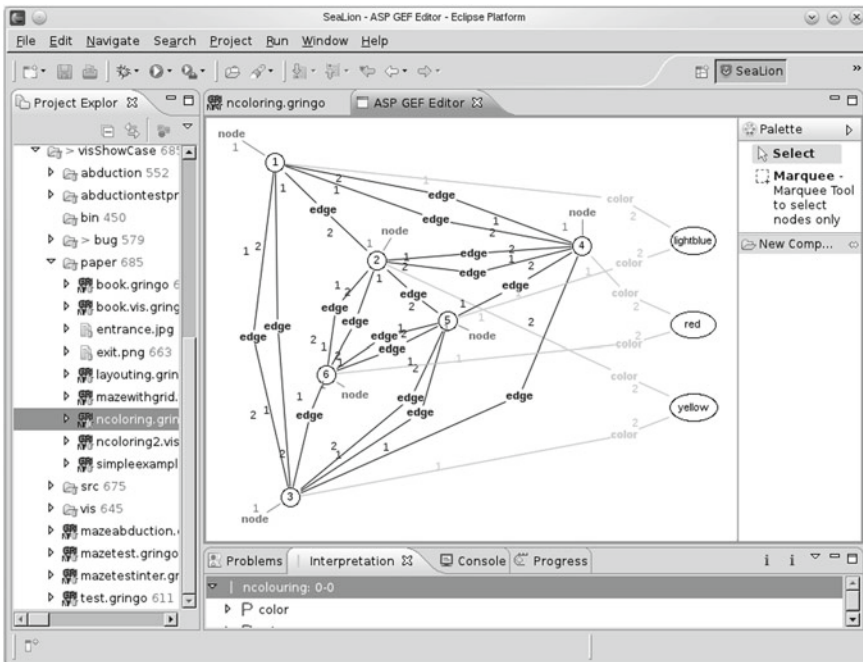


Fig. 11. SeaLion’s visual interpretation editor showing a generic visualisation of the graph colouring interpretation of Example 4 (the layout has been manually optimised).

interpretations should be visualised. Besides the features beyond visualisation, viz. the framework for editing visualisations and the support for multiple ASP solvers, there are also differences between **Kara** and these tools regarding visualisation aspects.

Kara allows to write more high-level specifications for positioning the graphical elements of a visualisation. While **IDPDraw** and **ASPVIZ** require the use of absolute coordinates, **Kara** additionally supports relative positioning and automatic layouting for graph and grid structures. Note that, technically, the former is realised by using ASP for guessing positions of the individual elements and adding respective constraints to ensure the correct layout, while the latter is realised by using a standard graph layouting algorithm which is part of the Eclipse framework. In **Kara**, as well as in **IDPDraw**, each graphical element has a unique identifier that can be used, e.g., to link elements or to set their properties (e.g., colour or font style). That way, programs can be written in a clear and elegant way since not all properties of an element have to be specified within a single atom. Here, **Kara** exploits that the latest ASP solvers support function symbols that allow for generating new identifiers from terms of the interpretation to visualise. **IDPDraw** does not support function symbols, however. Instead, for compound identifiers, **IDPDraw** uses predicates of variable length (e.g., *idp_polygon(id₁, id₂, ...)*). A disadvantage of this approach is that some solvers, like DLV, do not support predicates of variable length. **ASPVIZ** does not support identifiers for graphical objects at all. However, it does allow for defining named properties used for drawing, e.g., brushes of different colours and sizes can be assigned to constants that are then needed to be set in the command for drawing lines.

The support for a *z*-axis to determine which object should be drawn over others is available in **Kara** and **IDPDraw** but missing in **ASPVIZ**. However, a new version of **ASPVIZ** supports visualisations by 3-dimensional objects. Both **Kara** and **ASPVIZ** support the export of visualisations as vector graphics in the SVG format, which is not possible with **IDPDraw**. A feature that is supported by **ASPVIZ** and **IDPDraw**, however, is creating animations which is not possible with **Kara** so far. Note that an export feature for animations could be easily integrated in **Kara** as well.

Kara and **ASPVIZ** are written in Java and depend only on a Java Virtual Machine. **IDPDraw**, on the other hand, is written in C++ and depends on the qt libraries. Finally, **Kara** is embedded in an IDE, whereas **ASPVIZ** and **IDPDraw** are stand-alone tools.

Recently, a further visualisation tool for answer-sets that uses ASP for specifying visualisations has been written [12]. This system, called **Lonsdaleite**, is a lightweight Python script for visualising graph structures by mapping the atoms in an answer set to the input format of the **graphviz** utilities [13]. Thus, it only supports rendering a problem in a graph, but not other elements available in **ASPVIZ**, **IDPDraw**, and **Kara**.

A related approach from software engineering is the **Alloy Analyzer**, a tool to support the analysis of declarative software models [14]. Models are formulated in a first-order based specification language. The user may define object signatures, which can have properties (e.g., relationships to other objects). The tool can find satisfying instances of a model using translations to SAT. Since the approach is based on finding models for declarative specifications, it can be regarded as a form of ASP in a broader sense. The derived model instances are first-order structures that can be automatically visualised as graphs, where objects are represented as nodes and their relationships are represented as edges. Hence, visualisation in **Alloy** is closely related to the generic visualisation mode of **Kara** where also no dedicated visualisation program is needed. Finally, a useful feature of **Alloy** is filtering predicates and arguments away of the graph.

6 Conclusion

We presented the tool **Kara** for visualising and visual editing of interpretations in ASP. It supports generic as well as customised visualisations. A powerful language for defining a visualisation by means of ASP is provided, supporting, e.g., graph layouting, grids of graphical elements, and relative positioning. The editing feature uses abductive reasoning, inferring a new interpretation as hypothesis to explain a modified visualisation.

In future work, we want to add support for defining input and output signatures for programs in **SeaLion**. Then, the abduction framework of **Kara** could be easily extended such that one can derive inputs for a domain program such that one of its answer sets corresponds to a modified visualisation. We also consider adding a feature similar to the filtering mode of **Alloy** for getting a clearer generic visualisation. Moreover, we want to investigate model-driven development for ASP involving domain models that allow for obtaining generic visualisations that take structural information into account.

A Predefined Visualisation Predicates in Kara

| Atom | Intended meaning |
|--|--|
| $visellipse(id, height, width)$ | Defines an ellipse with specified height and width. |
| $visrect(id, height, width)$ | Defines a rectangle with specified height and width. |
| $vispolygon(id, x, y, ord)$ | Defines a point of a polygon. The ordering defines in which order the defined points are connected with each other. |
| $visimage(id, path)$ | Defines an image given in the specified file. |
| $visline(id, x_1, y_1, x_2, y_2, z)$ | Defines a line between the points (x_1, y_1) and (x_2, y_2) . |
| $visgrid(id, rows, cols, height, width)$ | Defines a grid with the specified number of rows and columns; <i>height</i> and <i>width</i> determine the grid size. |
| $visgraph(id)$ | Defines a graph. |
| $vistext(id, text)$ | Defines a text element. |
| $vislabel(id_g, id_t)$ | Sets the text element id_t as a label for graphical element id_g . Labels are supported for the following elements: $visellipse/3$, $visrect/3$, $vispolygon/4$, and $visconnect/3$. |
| $visinode(id_n, id_g)$ | Adds the graphical element id_n as a node to a graph id_g for automatic layouting. The following elements are supported as nodes: $visrect/3$, $visellipse/3$, $vispolygon/4$, $visimage/2$. |
| $visscale(id, height, weight)$ | Scales an image to the specified height and width. |
| $visposition(id, x, y, z)$ | Puts an element id on the fixed position (x, y, z) . |
| $visfontfamily(id, ff)$ | Sets the specified font ff for a text element id . |
| $visfontsize(id, size)$ | Sets the font size $size$ for a text element id . |
| $visfontstyle(id, style)$ | Sets the font style for a text element id to bold or italics. |
| $viscolor(id, color)$ | Sets the foreground colour for the element id . |
| $visbackgroundcolor(id, color)$ | Sets the background colour for the element id . |
| $visfillgrid(id_g, id_c, row, col)$ | Puts element id_c in cell (row, col) of the grid id_g . |
| $visconnect(id_c, id_{g_1}, id_{g_2})$ | Connects two elements, id_{g_1} and id_{g_2} , by a line such that id_{g_1} is the source and id_{g_2} is the target of the connection. |
| $visourcedeco(id, deco)$ | Sets the source decoration for a connection. |
| $vistargetdeco(id, deco)$ | Sets the target decoration for a connection. |
| $visleft(id_l, id_r)$ | Ensures that the x -coordinate of id_l is less than that of id_r . |
| $visright(id_r, id_l)$ | Ensures that the x -coordinate of id_r is greater than that of id_l . |
| $visabove(id_t, id_b)$ | Ensures that the y -coordinate of id_t is smaller than that of id_b . |
| $visbelow(id_b, id_t)$ | Ensures that the y -coordinate of id_b is greater than that of id_t . |
| $visinfrontof(id_1, id_2)$ | Ensures that the z -coordinate of id_1 is greater than that of id_2 . |
| $vishide(id)$ | Hides the element id . |
| $visdeletable(id)$ | Defines that the element id can be deleted in the visual editor. |
| $viscreatable(id)$ | Defines that the element id can be created in the visual editor. |
| $vischangable(id, prop)$ | Defines that the property $prop$ can be changed for the element id in the visual editor. |
| $vispossiblegridvalues(id, id_e)$ | Defines that graphical element id_e is available as possible grid value for a grid id in the visual editor. |

References

1. Baral, C.: Knowledge Representation, Reasoning, and Declarative Problem Solving. Cambridge University Press, Cambridge (2003)
2. Shapiro, E.Y.: Algorithmic program debugging. Ph.D. thesis, Yale University, New Haven, CT, USA (1982)
3. Oetsch, J., Pührer, J., Tompits, H.: Catching the Ouroboros: on debugging non-ground answer-set programs. *Theor. Pract. Logic Program.* **10**(4–5), 513–529 (2010)
4. Janhunen, T., Oikarinen, E., Tompits, H., Woltran, S.: Modularity aspects of disjunctive stable models. *J Artif. Intell. Res.* **35**, 813–857 (2009)
5. Cliffe, O., De Vos, M., Brain, M., Padget, J.: ASPVIZ: declarative visualisation and animation using answer set programming. In: Garcia de la Banda, M., Pontelli, E. (eds.) *ICLP 2008*. LNCS, vol. 5366, pp. 724–728. Springer, Heidelberg (2008)
6. Wittocx, J.: KRR Software: IDPDraw. <https://dtai.cs.kuleuven.be/krr/software/visualisation>
7. Peirce, C.S.: Abduction and induction. In: Buchler, J. (ed.) *Philosophical Writings of C.S Peirce*, Chapter 11, pp. 150–156. Dover, New York (1955)
8. Oetsch, J., Pührer, J., Tompits, H.: The SeaLion has landed: An IDE for answer-set programming—preliminary report. In: Tompits, H., Abreu, S., Oetsch, J., Pührer, J., Seipel, D., Umeda, M., Wolf, A. (eds.) *INAP/WLP 2011*. LNCS, vol. 7773, pp. 305–324. Springer, Heidelberg (2013)
9. Oetsch, J., Pührer, J., Tompits, H.: Methods and methodologies for developing answer-set programs - Project description. In: *Leibniz International Proceedings in Informatics (LIPIcs) of Technical Communications of the 26th International Conference on Logic Programming (ICLP 2010)*, Dagstuhl, Germany, Schloss Dagstuhl-Leibniz-Zentrum für Informatik, vol. 7, pp. 154–161 (2010)
10. The Eclipse Foundation: GEF (Graphical Editing Framework). <http://www.eclipse.org/gef/>
11. Denecker, M., Vennekens, J., Bond, S., Gebser, M., Truszczyński, M.: The second answer set programming competition. In: Erdem, E., Lin, F., Schaub, T. (eds.) *LPNMR 2009*. LNCS, vol. 5753, pp. 637–654. Springer, Heidelberg (2009)
12. Smith, A.: Lonsdaleite. <https://github.com/rndmcnll/Lonsdaleite>
13. AT&T Labs Research and Contributors: Graphviz. <http://www.graphviz.org/>
14. Jackson, D.: *Software Abstractions—Logic, Language, and Analysis*. MIT Press, Cambridge (2006)