

Deeply Semantic Inductive Spatio-Temporal Learning

Jakob Suchan¹, Mehul Bhatt¹, and Carl Schultz²

Human-Centred Cognitive Assistance¹ hcc.uni-bremen.de and
The DesignSpace Group,^{1,2} www.design-space.org
University of Bremen¹, and University of Münster², GERMANY

Abstract. We present an inductive spatio-temporal learning framework rooted in inductive logic programming. With an emphasis on visuo-spatial language, logic, and cognition, the framework supports learning with relational spatio-temporal features identifiable in a range of domains involving the processing and interpretation of dynamic visuo-spatial imagery. We present a prototypical system, and an example application in the domain of computing for visual arts and computational cognitive science.

Keywords: Spatio-Temporal Learning; Dynamic Visuo-Spatial Imagery; Declarative Spatial Reasoning; Inductive Logic Programming; AI and Art

1 INTRODUCTION

Cognitive assistive technologies and computer-human interaction systems involving an interplay of space, dynamics, and cognition necessitate capabilities for explainable reasoning, learning, and control about *space*, *actions*, *change*, and *interaction* [1]. Prime application scenarios, for instance, include (A1–A5): **(A1)**, activity grounding from video and point-clouds; **(A2)**, modelling and analysis of environmental processes at the geospatial scale; **(A3)**, medical computing scenarios replete with visuo-spatial imagery; **(A4)**, visuo-locomotive human behavioural data concerning aspects such as mobility or navigation, eye-tracking based visual perception research; **(A5)**, embodied human-machine interaction and control for commonsense cognitive robotics. A crucial requirement in relevant application contexts (such as A1–A5) pertains to the semantic interpretation of multi-modal human behavioural or socio-environmental data, with objectives ranging from knowledge acquisition (e.g., medical computing, computer-aided learning) and data analyses (e.g., activity interpretation) to hypothesis formation in experimental settings (e.g., empirical visual perception studies). The focus of our research is the processing and interpretation of dynamic visuo-spatial imagery with a particular emphasis on the ability to learn commonsense knowledge that is semantically founded in spatial, temporal, and spatio-temporal relations and patterns.

DEEP VISUO-SPATIAL SEMANTICS The high-level semantic interpretation and qualitative analysis of dynamic visuo-spatial imagery requires the representational and inferential mediation of commonsense abstractions of *space*, *time*, *action*, *change*, *interaction* and their mutual interplay thereof. In this backdrop, *deep visuo-spatial semantics* denotes the existence of declaratively grounded models —e.g., pertaining to

space, time, space-time, motion, actions & events, spatio-linguistic conceptual knowledge— and systematic formalisation supporting capabilities such as: (a). mixed quantitative qualitative spatial inference and question answering (e.g., about consistency, qualification and quantification of relational knowledge); (b). non-monotonic spatial reasoning (e.g., for abductive explanation); (c). relational learning of spatio-temporally grounded concepts; (d). integrated inductive-abductive spatio-temporal inference; (e). probabilistic spatio-temporal inference; (f). embodied grounding and simulation from the viewpoint of cognitive linguistics (e.g., for knowledge acquisition and inference based on natural language).

Recent perspectives on deep visuo-spatial semantics encompass methods for declarative (spatial) representation and reasoning —e.g., about *space and motion*— within frameworks such as constraint logic programming (rule-based spatio-temporal inference [4, 24]), answer-set programming (for non-monotonic spatial reasoning [27]), description logics (for spatio-terminological reasoning [3]), inductive logic programming (for inductive-abductive spatio-temporal learning [5, 6]) and other specialised forms of commonsense reasoning based on expressive action description languages for modelling *space, events, action, and change* [1, 2]. In general, deep visuo-spatial semantics driven by declarative spatial representation and reasoning pertaining to dynamic visuo-spatial imagery is relevant and applicable in a variety of cognitive interaction systems and assistive technologies at the interface of (spatial) language, (spatial) logic, and (visuo-spatial) cognition.

INDUCTIVE SPATIO-TEMPORAL LEARNING (WITH DEEP SEMANTICS)

This research is motivated by the need to have a systematic inductive logic programming [15] founded spatio-temporal learning framework and corresponding system that:

- provides an expressive spatio-linguistically motivated ontology to predicate primitive and complex (domain-independent) relational spatio-temporal features identifiable in a broad range of application domains (e.g., A1–A5) involving the processing and interpretation of dynamic visuo-spatial imagery.
- supports spatio-temporal relations natively such that the semantics of these relations is directly built into the underlying ILP-based learning framework.
- supports seamless mixing of, and transition between, quantitative and qualitative spatial data.

We particularly emphasise and ensure compatibility with the general setup of (constraint) logic programming framework such that diverse knowledge sources and reasoning mechanisms outside of inductive learning may be directly interfaced, and reasoning / learning capabilities be combined within large-scale integrated systems for cognitive computing.

2 LEARNING FROM RELATIONAL SPATIO-TEMPORAL STRUCTURE: A GENERAL FRAMEWORK AND SYSTEM

We present a general framework and working prototype for an inductive spatio-temporal learning system with an elaborate ontology supporting a range of space-time features;

we demonstrate the functional capabilities from the viewpoint of AI-based computing for the arts & social sciences, and computational cognitive science.

2.1 THE SPATIO-TEMPORAL DOMAIN \mathcal{O}_{SP} , AND \mathcal{QS}

The spatio-temporal ontology $\mathcal{O}_{sp} \equiv_{def} \langle \mathcal{E}, \mathcal{R} \rangle$ is characterised by the basic spatial entities (\mathcal{E}) that can be used as abstract representations of domain-objects and the relational spatio-temporal structure (\mathcal{R}) that characterises the qualitative spatio-temporal relationships amongst the supported entities in (\mathcal{E}). The following primitive spatial entities are sufficient to characterise the learning mechanism and its sample application for this paper:

a **point** is a pair of reals x, y ; a **vector** is a pair of reals v_x, v_y ; an **oriented point** consists of a point p and a vector v ; a **line segment** is a pair of end points p_1, p_2 ($p_1 \neq p_2$); a **rectangle** is a point p representing the bottom left corner, a direction vector v defining the orientation of the base of the rectangle, and a real width and height w, h ($0 < w, 0 < h$); an **axis-aligned rectangle** is a rectangle with fixed direction vector $v = (1, 0)$; a **circle** is a centre point p and a real radius r ($0 < r$); a **simple polygon** is defined by a list of n vertices (points) p_1, \dots, p_n (spatially ordered counter-clockwise) such that the boundary is non-self-intersecting, i.e., there does not exist a polygon boundary edge between vertices p_i, p_{i+1} that intersects some other edge p_j, p_{j+1} for all $1 \leq i < j < n$ and $i + 1 < j$.

Spatio-temporal relationships (\mathcal{R}) between the basic entities in \mathcal{E} may be characterised with respect to arbitrary spatial and spatio-temporal domains such as *mereotopology*, *orientation*, *distance*, *size*, *motion*; Table 1 lists the relevant supported relations from the viewpoint of established spatial abstraction calculi such as the Region Connection Calculus [16], Rectangle Algebra and Block Algebra [7], LR Calculus [20], Oriented-Point Relation Algebra (OPRA) [14], and Space-Time Histories [8, 9].

\mathcal{QS} – ANALYTIC SEMANTICS FOR \mathcal{O}_{SP} We adopt an analytic approach to spatial reasoning, where the semantics of spatial relations are encoded as polynomial constraints within a (constraint) logic programming setup. The analytic method supports the integration of qualitative and quantitative spatial information, and provides a means for sound, complete and approximate spatial reasoning [4]. For example, let axis-aligned rectangles a, b each be defined by a bottom-left vertex (x_i, y_i) and a width and height w_i, h_i , for $i \in \{a, b\}$ such that x_i, y_i, w_i, h_i are reals. The relation that a is a *non-tangential proper part* of b corresponds to the polynomial constraint:

$$(x_b < x_a) \wedge (x_a + w_a < x_b + w_b) \wedge (y_b < y_a) \wedge (y_a + h_a < y_b + h_b)$$

Continuing with the example, this is generalised to arbitrarily oriented rectangles. Determining whether a point is inside an arbitrary rectangle is based on vector projection. Point p is projected onto vector v by taking the dot product:

$$(x_p, y_p) \cdot (x_v, y_v) = x_p x_v + y_p y_v.$$

With this approach, the task of determining whether a set of spatial relations is consistent then becomes the task of determining whether a system of polynomial constraints is satisfiable. We emphasise that our approach and framework are not limited to the above

SPATIAL DOMAIN (\mathcal{QS})	Formalisms	Spatial Relations (\mathcal{R})	Entities (\mathcal{E})
Mereotopology	RCC-5, RCC-8 [16]	disconnected (dc), external contact (ec), partial overlap (po), tangential proper part (tpp), non-tangential proper part (ntpp), proper part (pp), part of (p), discrete (dr), overlap (o), contact (c)	arbitrary rectangles, circles, polygons, cuboids, spheres
	Rectangle & Block algebra [7]	proceeds, meets, overlaps, starts, during, finishes, equals	axis-aligned rectangles and cuboids
Orientation	LR [20]	left, right, collinear, front, back, on	2D point, circle, polygon with 2D line
	OPRA [14]	facing towards, facing away, same direction, opposite direction	oriented points, 2D/3D vectors
Distance, Size	QDC [10]	adjacent, near, far, smaller, equi-sized, larger	rectangles, circles, polygons, cuboids, spheres
Dynamics, Motion	Space-Time Histories [8, 9]	moving: towards, away, parallel; growing / shrinking: vertically, horizontally; passing: in front, behind; splitting / merging	rectangles, circles, polygons, cuboids, spheres

Table 1. The Spatio-Temporal Domain \mathcal{O}_{sp} supported within the Learning Framework

entities; a wider class of 2D and 3D spatial entities are supported and may be defined as per domain-specific and computational needs [4, 18, 27, 19].

INDUCTIVE LEARNING WITH THE SPATIAL SYSTEM $\langle \mathcal{O}_{SP}, \mathcal{QS} \rangle$ Learning is founded on the Aleph ILP system [21]. Learning spatio-temporal structures, is based on integrating the spatial ontology \mathcal{O}_{sp} described above, into the basic learning setup of ILP.

Given: (1) A set of examples E , consisting of positive and negative examples for the desired spatio-temporal structure, i.e., $E = E^+ \cup E^-$, where each example is given by a set of spatio-temporal observations in the domain; (2) the (spatio-temporal) background knowledge B .

The *spatio-temporal learning domain* is defined by basic spatial entities (\mathcal{E}) constituting the domain objects, the relational spatial structure (\mathcal{R}) describing the spatio-temporal configuration of spatial entities in the domain, and rules defining spatio-temporal phenomena and characteristics of the domain. In this context, spatio-temporal facts characterising the learning examples E can be given as, (a) numerical representation of domain objects, (b) qualitative relations between spatial entities, or (c) a mixed qualitative-quantitative representations, where the facts are partially grounded in numerical observations.

Learning: The learning task is defined as finding hypothesis H consisting of spatio-temporal relations (\mathcal{R}) holding between basic spatial entities (\mathcal{E}), such that $H \cup B \models E^+$, and $H \cup B \not\models E^-$.

As such, the spatial ontology \mathcal{O}_{sp} constitutes an integrated part of the learning setup and spatio-temporal semantics are available throughout the learning process.

3 LEARNING CINEMATOGRAPHIC PATTERNS AND THEIR VISUAL RECEPTION: THE CASE OF SYMMETRY

Aimed at cognitive film studies and visual perception research, we present a use-case pertaining to the (visual) learning of cinematographic patterns of symmetry and its vi-




Fig. 1. Positive examples for symmetric scene structures at the object level

sual reception (by means of eye-tracking) by subjects.¹ To demonstrate the temporal aspect of the learning framework, we demonstrate the capability to learn “*axioms of visual perception*” from dynamic eye-tracking data; both the chosen films and their corresponding eye-tracking data are obtained from a large-scale experiment in visual perception of films [23, 22]. The presented example translates to a variety of cases involving visual perception and human behaviour studies.

Learning Spatial Structures: Object-Level Symmetry As an example for learning spatial structures, we consider symmetry in the relative object placement in a movie scene (see Fig. 1). In particular, learning is based on the spatial configuration of *people, faces*, and their *facing direction*, directly obtained from computer vision algorithms as described in [23]. In this context, *positive and negative examples*, are given as numerical spatial facts about domain objects in the image.

```
...
detection(id(0), image(3), class(person), rectangle(point(319, 194), 319, 456)).
detection(id(1), image(3), class(person), rectangle(point(678, 215), 367, 452)).
detection(id(0), image(3), class(face), rectangle(point(438, 246), 86, 86)).
detection(id(1), image(3), class(face), rectangle(point(745, 284), 87, 87)).
2d_facing_dir(id(0), image(3), vector(0.550864, 0.834595), magnitude(6.26042)).
2d_facing_dir(id(1), image(3), vector(-0.500519, 0.865726), magnitude(4.82556)).
...
```

We define representations of domain objects linking the numerical description of objects in the image to basic spatial entities describing different aspects of these objects, e.g. the bounding box (*rectangles*), or the center-point (*points*).

```
entity(center(person(P)), point(X, Y), image(Img)) :-
  detection(_, image(Img), class(person), rectangle(point(Xr, Yr), W, H)),
  X is Xr + W/2, Y is Yr + H/2.
```

In addition to the detected domain objects, we define abstract geometric objects needed to describe symmetry, e.g. the *symmetry axis* in the center of the image.

```
entity(symmetry_obj(center_axis), line(X, 0, X, Y), image(Img)) :-
  img(image(Img)), media_size(size(MediaWidth, MediaHeight), image(Img)),
  X is MediaWidth/2, Y is MediaHeight.
```

Learning: We learn the relational spatial structure consisting of qualitative spatial relationships characterising symmetry in the configuration of the spatial entities in the image, i.e. we consider relations of *topology, orientation, distance, and size*.

```
:- modeh(1, symmetric(+img)).           :- modeb(*, entity(#ent, -obj, +img)).
:- modeb(*, topology(rcc8(#rel, +obj, +obj)). :- modeb(*, distance(#rel, +obj, +obj)).
:- modeb(*, orientation(#rel, +obj, +obj)). :- modeb(*, size(#rel, +obj, +obj)).
```

¹ Our case-study is motivated by a broader multi-level interpretation of symmetry from the viewpoint of film cinematography [25]; however, the specific example of this paper focusses on one aspect of this multi-level symmetry characterisation involving relative object placement in a movie scene.

Exemplary symmetrical spatial structures, learned by the system include the following.

```
symmetric(A) :- entity(center(person(0)),B,A), entity(center(person(1)),C,A),
               entity(symmetry_object(center_axis),D,A), distance(equidistant,D,C,B).

symmetric(A) :- entity(person(0),B,A), entity(person(1),C,A), size(same,C,B).
```

Learning Spatio-Temporal Dynamics: Axioms of Perception We illustrate learning of spatio-temporal dynamics in the context of visual perception, by learning perceptual patterns from eye-tracking data and people tracks in a movie scene. As an example we focus on attention of a person switching from one individual to another.

```
detection(id(0), frame(426), class(person), rectangle(point(385,66),244,271)).
detection(id(1), frame(426), class(person), rectangle(point(111,68),332,276)).
gaze(point(frame(426), point(859,212))).
```

Learning: We adapt the general learning setup of the example above, for learning spatio-temporal dynamics by introducing the predicate `holds_in/2` to denote that a spatial relation holds between two entities at a time point.

```
...
:- modeb(*, holds_in(topology(#rel, +ent, +ent), +time)).
:- modeb(*, time(#rel, +time, +time)).
...
```

Spatio-temporal dynamics constituting attention switches include the following.

```
att_switch(B) :- holds_in(topology(inside, gaze, person(1)), A),
                 holds_in(topology(inside, gaze, person(2)), B), time(consecutive, A, B).
```

4 DISCUSSION AND OUTLOOK

Directly comparable to this research is the line of work on integrated inductive-abductive reasoning for learning spatio-temporal relational models from video in [5, 6]; here, spatio-temporal learning in the context of ILP has only been addressed for the case of topological relations. Furthermore, the ILP learning framework does not have built-in semantics for the topological relations. Aside from this, learning relational spatial structures was investigated in the context of learning spatial relations from language [12], and within the geospatial domain [13, 26]. Probabilistic Logic Programming frameworks such as PRISM [17] and ProbLog [11] have been used for learning parameters, and the structure, of probabilistic logic programs, although (qualitative) spatial reasoning has not been directly addressed. The main point-of-departure of this paper with respect to the state of the art in (qualitative) spatial learning is that the semantics of spatial, temporal, and spatio-temporal relations are directly built within the inductive learning framework of ILP. Pragmatically, what this implies is that it is possible to seamlessly describe a learning problem using a generic relational spatio-temporal ontology directly as part of a logic programming based learning environment. To the best of our knowledge, such a general spatio-temporal learning framework with built in semantics for mixed qualitative-quantitative spatio-temporal reasoning capabilities has not been available before. Furthermore, the ontology of space-time features supported in our framework goes much beyond topological relations addressing orientation, distance, and size. Future research will focus on enhancing the expressivity of the spatio-temporal relations to cover a wider range of domain-independent features characterising spatio-temporal dynamics.

Bibliography

- [1] M. Bhatt. Reasoning about Space, Actions and Change: A Paradigm for Applications of Spatial Reasoning. In *Qualitative Spatial Representation and Reasoning: Trends and Future Directions*. IGI Global, USA, 2012. ISBN ISBN13: 9781616928681.
- [2] M. Bhatt and S. Loke. Modelling Dynamic Spatial Systems in the Situation Calculus. *Spatial Cognition and Computation*, 8(1):86–130, 2008. ISSN 1387-5868.
- [3] M. Bhatt, F. Dylla, and J. Hois. Spatio-terminological inference for the design of ambient environments. In *Spatial Information Theory, 9th International Conference, COSIT 2009, Aber Wrac’h, France, September 21-25, 2009, Proceedings*, volume 5756 of *Lecture Notes in Computer Science*, pages 371–391. Springer, 2009. ISBN 978-3-642-03831-0. doi: 10.1007/978-3-642-03832-7_23.
- [4] M. Bhatt, J. H. Lee, and C. P. L. Schultze. CLP(QS): A declarative spatial reasoning framework. In *Spatial Information Theory - 10th International Conference, COSIT 2011, Belfast, ME, USA, September 12-16, 2011. Proceedings*, pages 210–230, 2011. doi: 10.1007/978-3-642-23196-4_12.
- [5] K. S. R. Dubba, M. Bhatt, F. Dylla, D. C. Hogg, and A. G. Cohn. Interleaved inductive-abductive reasoning for learning complex event models. In S. Muggleton, A. Tamaddoni-Nezhad, and F. A. Lisi, editors, *Inductive Logic Programming - 21st International Conference, ILP 2011, Windsor Great Park, UK, July 31 - August 3, 2011, Revised Selected Papers*, volume 7207 of *Lecture Notes in Computer Science*, pages 113–129. Springer, 2011. ISBN 978-3-642-31950-1. doi: 10.1007/978-3-642-31951-8_14.
- [6] K. S. R. Dubba, A. G. Cohn, D. C. Hogg, M. Bhatt, and F. Dylla. Learning relational event models from video. *J. Artif. Intell. Res. (JAIR)*, 53:41–90, 2015. doi: 10.1613/jair.4395.
- [7] H. W. Guesgen. *Spatial reasoning based on Allen’s temporal logic*. Technical Report TR-89-049. International Computer Science Institute Berkeley, 1989.
- [8] P. J. Hayes. Naive physics I: ontology for liquids. In J. R. Hubbs and R. C. Moore, editors, *Formal Theories of the Commonsense World*. Ablex Publishing Corporation, Norwood, NJ, 1985.
- [9] S. M. Hazarika. *Qualitative spatial change: space-time histories and continuity*. PhD thesis, The University of Leeds, 2005.
- [10] D. Hernández, E. Clementini, and P. Di Felice. *Qualitative distances*. Springer, 1995.
- [11] A. Kimmig, L. De Raedt, and H. Toivonen. Probabilistic explanation based learning. In *European Conference on Machine Learning*, pages 176–187. Springer, 2007.
- [12] P. Kordjamshidi, P. Frasconi, M. van Otterlo, M. Moens, and L. D. Raedt. Relational learning for spatial relation extraction from natural language. In S. Muggleton, A. Tamaddoni-Nezhad, and F. A. Lisi, editors, *Inductive Logic Programming - 21st International Conference, ILP 2011, Windsor Great Park, UK, July 31 - August 3, 2011, Revised Selected Papers*, volume 7207 of *Lecture Notes in Computer Science*, pages 204–220. Springer, 2011. doi: 10.1007/978-3-642-31951-8_20.
- [13] D. Malerba and F. A. Lisi. Discovering associations between spatial objects: An ilp application. In *Proceedings of the 11th International Conference on Inductive Logic Programming, ILP ’01*, pages 156–163, London, UK, UK, 2001. Springer-Verlag. ISBN 3-540-42538-1.
- [14] R. Moratz. Representing relative direction as a binary relation of oriented points. In *ECAI*, pages 407–411, 2006.
- [15] S. Muggleton and L. D. Raedt. Inductive logic programming: Theory and methods. *JOURNAL OF LOGIC PROGRAMMING*, 19(20):629–679, 1994.

- [16] D. A. Randell, Z. Cui, and A. G. Cohn. A spatial logic based on regions and connection. *KR*, 92:165–176, 1992.
- [17] T. Sato and Y. Kameya. New advances in logic-based probabilistic modeling by prism. In *Probabilistic inductive logic programming*, pages 118–155. Springer, 2008.
- [18] C. Schultz and M. Bhatt. A numerical optimisation based characterisation of spatial reasoning. In *Rule Technologies. Research, Tools, and Applications: 10th International Symposium, RuleML 2016, Stony Brook, NY, USA, July 6-9, 2016. Proceedings*, pages 199–207. Springer International Publishing, 2016. ISBN 978-3-319-42019-6.
- [19] C. Schultz, M. Bhatt, and J. Suchan. Probabilistic spatial reasoning in constraint logic programming. In *Tenth International Conference on Scalable Uncertainty Management (SUM 2016) (to appear)*, 2016.
- [20] A. Scivos and B. Nebel. The Finest of its Class: The Natural, Point-Based Ternary Calculus LR for Qualitative Spatial Reasoning. In *C. Freksa et al. (2005), Spatial Cognition IV. Reasoning, Action, Interaction: International Conference Spatial Cognition. Lecture Notes in Computer Science Vol. 3343, Springer, Berlin Heidelberg*, volume 3343, pages 283–303, 2004.
- [21] A. Srinivasan. *The Aleph Manual*, 2001. URL <http://web.comlab.ox.ac.uk/oucl/research/areas/machlearn/Aleph/>.
- [22] J. Suchan and M. Bhatt. Semantic question-answering with video and eye-tracking data – ai foundations for human visual perception driven cognitive film studies. In *IJCAI 2016: 25th International Joint Conference on Artificial Intelligence*, New York City, USA, July 2016.
- [23] J. Suchan and M. Bhatt. The geometry of a scene: On deep semantics for visual perception driven cognitive film, studies. In *2016 IEEE Winter Conference on Applications of Computer Vision, WACV 2016, Lake Placid, NY, USA, March 7-10, 2016*, pages 1–9. IEEE Computer Society, 2016. ISBN 978-1-5090-0641-0. doi: 10.1109/WACV.2016.7477712.
- [24] J. Suchan, M. Bhatt, and P. E. Santos. Perceptual narratives of space and motion for semantic interpretation of visual data. In *Computer Vision - ECCV 2014 Workshops - Zurich, Switzerland, September 6-7 and 12, 2014, Proceedings, Part II*, pages 339–354, 2014. doi: 10.1007/978-3-319-16181-5_24.
- [25] J. Suchan, M. Bhatt, and S. Yu. The perception of symmetry in the moving image: multi-level computational analysis of cinematographic scene structure and its visual reception. In E. Jain and S. Jörg, editors, *ACM Symposium on Applied Perception, SAP 2016, Anaheim, California, USA, July 22-23, 2016*, page 142. ACM, 2016. ISBN 978-1-4503-4383-1. doi: 10.1145/2931002.2948721.
- [26] D. Vaz, V. S. Costa, and M. Ferreira. Fire! firing inductive rules from economic geography for fire risk detection. In P. Frasconi and F. A. Lisi, editors, *Inductive Logic Programming - 20th International Conference, ILP 2010, Florence, Italy, June 27-30, 2010. Revised Papers*, volume 6489 of *Lecture Notes in Computer Science*, pages 238–252. Springer, 2010. doi: 10.1007/978-3-642-21295-6_27.
- [27] P. A. Walega, M. Bhatt, and C. P. L. Schultz. ASPMT(QS): Non-Monotonic Spatial Reasoning with Answer Set Programming Modulo Theories. In *Logic Programming and Nonmonotonic Reasoning - 13th International Conference, LPNMR 2015, Lexington, KY, USA, September 27-30, 2015. Proceedings*, pages 488–501, 2015. doi: 10.1007/978-3-319-23264-5_41.