

Out of Sight But Not Out of Mind: An Answer Set Programming Based Online Abduction Framework for Visual Sensemaking in Autonomous Driving

Jakob Suchan^{1,3}, Mehul Bhatt^{2,3} and Srikrishna Varadarajan³

¹University of Bremen, Germany

²Örebro University, Sweden

³CoDesign Lab / Cognitive Vision

jsuchan@uni-bremen.de, mehul.bhatt@oru.se, krish@codesign-lab.org

Abstract

We demonstrate the need and potential of systematically integrated *vision* and *semantics* solutions for visual sensemaking (in the backdrop of autonomous driving). A general method for *online* visual sensemaking using answer set programming is systematically formalised and fully implemented. The method integrates state of the art in visual computing, and is developed as a modular framework usable within hybrid architectures for perception & control. We evaluate and demo with community established benchmarks KITTIMOD and MOT. As use-case, we focus on the significance of human-centred visual sensemaking —e.g., semantic representation and explainability, question-answering, commonsense interpolation— in safety-critical autonomous driving situations.

1 Motivation

Autonomous driving research has received enormous academic & industrial interest in recent years. This surge has coincided with (and been driven by) advances in *deep learning* based computer vision research. Although deep learning based vision & control has (arguably) been successful for self-driving vehicles, we posit that there is a clear need and tremendous potential for hybrid visual sensemaking solutions (integrating *vision and semantics*) towards fulfilling essential legal and ethical responsibilities involving explainability, human-centred AI, and industrial standardisation (e.g., pertaining to representation, realisation of rules and norms).

Autonomous Driving: “Standardisation & Regulation”

As the self-driving vehicle industry develops, it will be necessary —e.g., similar to sectors such as medical computing, computer aided design— to have an articulation and community consensus on aspects such as representation, interoperability, human-centred performance benchmarks, and data archival & retrieval mechanisms.¹ In spite of major investments in self-driving vehicle research, issues related to

¹ Within autonomous driving, the need for standardisation and ethical regulation has most recently garnered interest internationally, e.g., with the Federal Ministry of Transport and Digital Infrastruc-




Figure 1: Out of sight but not out of mind; the case of hidden entities: an occluded cyclist.

human-centred’ness, human collaboration, and standardisation have been barely addressed, with the current focus in driving research primarily being on two basic considerations: *how fast to drive, and which way and how much to steer*. This is necessary, but inadequate if autonomous vehicles are to become commonplace and function with humans. Ethically driven standardisation and regulation will require addressing challenges in semantic visual interpretation, natural / multimodal human-machine interaction, high-level data analytics (e.g., for post hoc diagnostics, dispute settlement) etc. This will necessitate —amongst other things— human-centred qualitative benchmarks and multifaceted hybrid AI solutions.

Visual Sensemaking Needs Both “Vision & Semantics”

We demonstrate the significance of semantically-driven methods rooted in knowledge representation and reasoning (KR) in addressing research questions pertaining to explainability and human-centred AI particularly from the viewpoint of sensemaking of dynamic visual imagery. Consider the *occlusion scenario* in Fig. 1:

Car (c) is in-front; during this time, person (p) is on a bicycle (b) and positioned front-right of c and moving-forward. Car c turns-right, during which the bicyclist $\langle p, b \rangle$ is not visible. Subsequently, bicyclist $\langle p, b \rangle$ reappears.

The occlusion scenario indicates several challenges concerning aspects such as: identity maintenance, making default assumptions, computing counterfactuals, projection, and interpolation of missing information (e.g., what could be hypoth-

ture in Germany taking a lead in eliciting 20 key propositions (with legal implications) for the fulfilment of ethical commitments for automated and connected driving systems [BMVI, 2018].

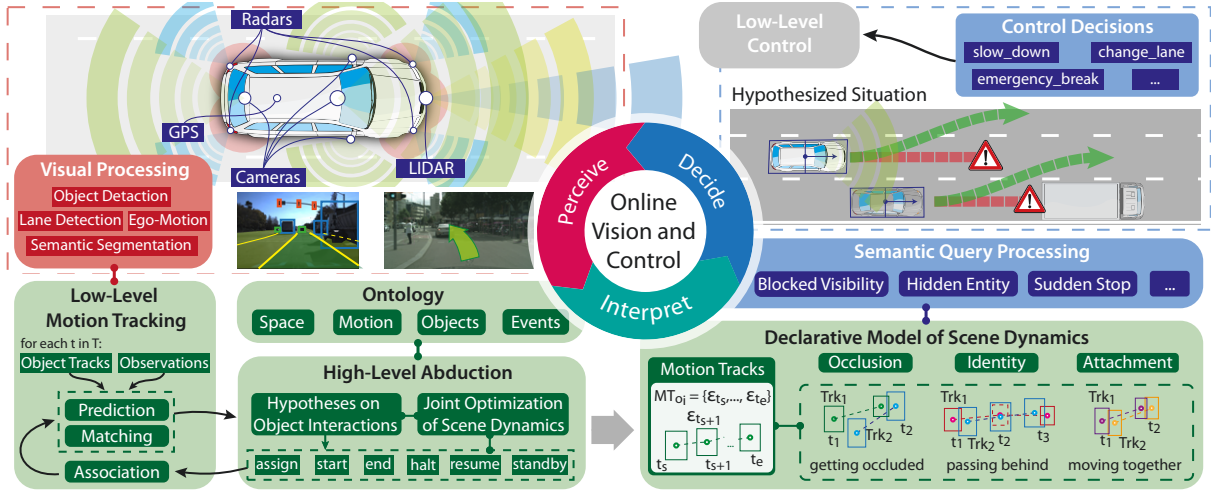


Figure 2: A General Online Abduction Framework / Conceptual Overview

esised about bicyclist $\langle p, b \rangle$ when it is *occluded*; how can this hypothesis enable in planning an immediate next step). Addressing such challenges —be it realtime or post-hoc— in view of human-centred AI concerns pertaining to ethics, explainability and regulation requires a systematic integration of **Semantics and Vision**, i.e., robust commonsense representation & inference about spacetime dynamics on the one hand, and powerful low-level visual computing capabilities, e.g., pertaining to object detection and tracking on the other.

Key Contributions. We develop a general and systematic declarative visual sensemaking method capable of online abduction: *realtime, incremental, commonsense* question-answering and belief maintenance over dynamic visuospatial imagery. Supported are (1–3): (1). human-centric representations semantically rooted in spatio-linguistic primitives as they occur in natural language [Bhatt *et al.*, 2013; Mani and Pustejovsky, 2012]; (2). driven by Answer Set Programming (ASP) [Brewka *et al.*, 2011], the ability to abductively compute commonsense interpretations and explanations in a range of (a)typical everyday driving situations, e.g., concerning safety-critical decision-making; (3). online performance of the overall framework modularly integrating high-level commonsense reasoning and state of the art low-level visual computing for practical application in real world settings. We present the formal framework & its implementation, and demo & empirically evaluate with community established real-world datasets and benchmarks, namely: KIT-TIMOD [Geiger *et al.*, 2012] and MOT [Milan *et al.*, 2016].

2 Visual Sensemaking: A General Method Driven by ASP

Our proposed framework, in essence, jointly solves the problem of assignment of *detections* to *tracks* and explaining overall scene dynamics (e.g. appearance, disappearance) in terms of high-level *events* within an online integrated low-level visual computing and high level abductive reasoning framework (Fig. 2). Rooted in answer set programming, the framework is general, modular, and designed for integration

as a reasoning engine within (hybrid) architectures designed for real-time decision-making and control where visual perception is needed as one of the several components. In such large scale AI systems the declarative model of the scene dynamics resulting from the presented framework can be used for semantic Q/A, inference etc. to support decision-making.

2.1 Space, Motion, Objects, Events

Reasoning about dynamics is based on high-level representations of objects and their respective motion & mutual interactions in spacetime. Ontological primitives for commonsense reasoning about spacetime (Σ_{st}) and dynamics (Σ_{dyn}) are:

Σ_{st} : **domain-objects** $\mathcal{O} = \{o_1, \dots, o_n\}$ represent the visual elements in the scene, e.g., *people, cars, cyclists*; elements in \mathcal{O} are geometrically interpreted as **spatial entities** $\mathcal{E} = \{\varepsilon_1, \dots, \varepsilon_n\}$; spatial entities \mathcal{E} may be regarded as *points, line-segments* or (axis-aligned) *rectangles* based on their spatial properties (and a particular reasoning task at hand). The temporal dimension is represented by **time points** $\mathcal{T} = \{t_1, \dots, t_n\}$. $\mathcal{MT}_{o_i} = (\varepsilon_{t_s}, \dots, \varepsilon_{t_e})$ represents the **motion track** of a single object o_i , where t_s and t_e denote the start and end time of the track and ε_{t_s} to ε_{t_e} denotes the spatial entity (\mathcal{E}) —e.g., the *axis-aligned bounding box*—corresponding to the object o_i at time points t_s to t_e . The spatial configuration of the scene and changes within it are characterised based on the qualitative **spatio-temporal relationships** (\mathcal{R}) between the domain objects. For the running and demo examples of this paper, positional relations on axis-aligned rectangles based on the rectangle algebra (RA) [Balbiani *et al.*, 1999] suffice; RA uses the relations of Interval Algebra (IA) [Allen, 1983] $\mathcal{R}_{IA} \equiv \{\text{before, after, during, contains, starts, started_by, finishes, finished_by, overlaps, overlapped_by, meets, met_by, equal}\}$ to relate two objects by the *interval relations* projected along each dimension separately (e.g., horizontal and vertical dimensions).

Σ_{dyn} : The set of **fluents** $\Phi = \{\phi_1, \dots, \phi_n\}$ and **events** $\Theta = \{\theta_1, \dots, \theta_n\}$ respectively characterise the dynamic properties of the objects in the scene and high-level abducibles (Table

Algorithm 1: Online_Abduction(\mathcal{V}, Σ)

Data: Visual imagery (\mathcal{V}), and background knowledge $\Sigma \equiv_{def} \Sigma_{dyn} \cup \Sigma_{st}$

Result: Visual Explanations ($\mathcal{E}\mathcal{X}\mathcal{P}$) (Refer Fig. 3)

```

1  $\mathcal{MT} \leftarrow \emptyset, \mathcal{H}^{events} \leftarrow \emptyset$ 
2 for  $t \in T$  do
3    $\mathcal{VO}_t \leftarrow observe(\mathcal{V}_t)$ 
4    $\mathcal{P}_t \leftarrow \emptyset, \mathcal{ML}_t \leftarrow \emptyset$ 
5   for  $trk \in \mathcal{MT}_{t-1}$  do
6      $p_{trk} \leftarrow kalman.predict(trk)$ 
7      $\mathcal{P}_t \leftarrow \mathcal{P}_t \cup p_{trk}$ 
8     for  $obs \in \mathcal{VO}_t$  do
9        $ml_{trk,obs} \leftarrow calc\_IoU(p_{trk}, obs)$ 
10       $\mathcal{ML}_t \leftarrow \mathcal{ML}_t \cup ml_{trk,obs}$ 
11    $abduce(\langle \mathcal{H}_t^{assign}, \mathcal{H}_t^{events} \rangle)$  (Refer Step 2)
12    $\mathcal{H}^{events} \leftarrow \mathcal{H}^{events} \cup \mathcal{H}_t^{events}$ 
13    $\mathcal{MT}_t \leftarrow update(\mathcal{MT}_{t-1}, \mathcal{VO}_t, \mathcal{H}_{assign})$ 
14 return  $\mathcal{E}\mathcal{X}\mathcal{P} \leftarrow \langle \mathcal{H}^{events}, \mathcal{MT} \rangle$ 

```

1). For reasoning about dynamics (with $\langle \Phi, \Theta \rangle$), we use a variant of event calculus as per [Ma *et al.*, 2014; Miller *et al.*, 2013]; in particular, for examples of this paper, the functional event calculus fragment (Σ_{dyn}) of Ma *et al.* [2014] suffices: main axioms relevant here pertain to $occurs_at(\theta, t)$ denoting that an event occurred at time t and $holds_at(\phi, v, t)$ denoting that v holds for a fluent ϕ at time t .²

$$\Sigma \equiv_{def} \Sigma_{dyn} \langle \Phi, \Theta \rangle \cup \Sigma_{st} \langle \mathcal{O}, \mathcal{E}, \mathcal{T}, \mathcal{MT}, \mathcal{R} \rangle$$

2.2 Tracking as Abduction

Scene dynamics are tracked using a *detect and track* approach: we tightly integrate low-level visual computing (for detecting scene elements) with high-level ASP-based abduction to solve the assignment of observations to object tracks in an *incremental* manner. For each time point t we generate a *problem specification* consisting of the object tracks and visual observations and use ASP to abductively solve the corresponding assignment problem incorporating the ontological structure of the domain / data (abstracted with Σ). **Steps 1–3** (Alg. 1 & Fig. 3) are as follows:

Step 1. Formulating the Problem Specification The ASP problem specification for each time point t is given by the tuple $\langle \mathcal{VO}_t, \mathcal{P}_t, \mathcal{ML}_t \rangle$ and the sequence of events (\mathcal{H}^{events}) before time point t .

Visual Observations (\mathcal{VO}_t): Scene elements derived directly from the visual input data are represented as spatial entities \mathcal{E} , i.e., $\mathcal{VO}_t = \{\varepsilon_{obs_1}, \dots, \varepsilon_{obs_n}\}$ is the set of observations at time t (Fig. 3). For the examples and empirical evaluation in this paper (Sec. 3) we focus on *Obstacle / Object Detections* – detecting cars, pedestrians, cyclists, traffic lights

²ASP encoding of the domain independent axioms of the Functional Event Calculus (FEC) used as per: <https://www.ucl.ac.uk/infostudies/efec/fec.lp>

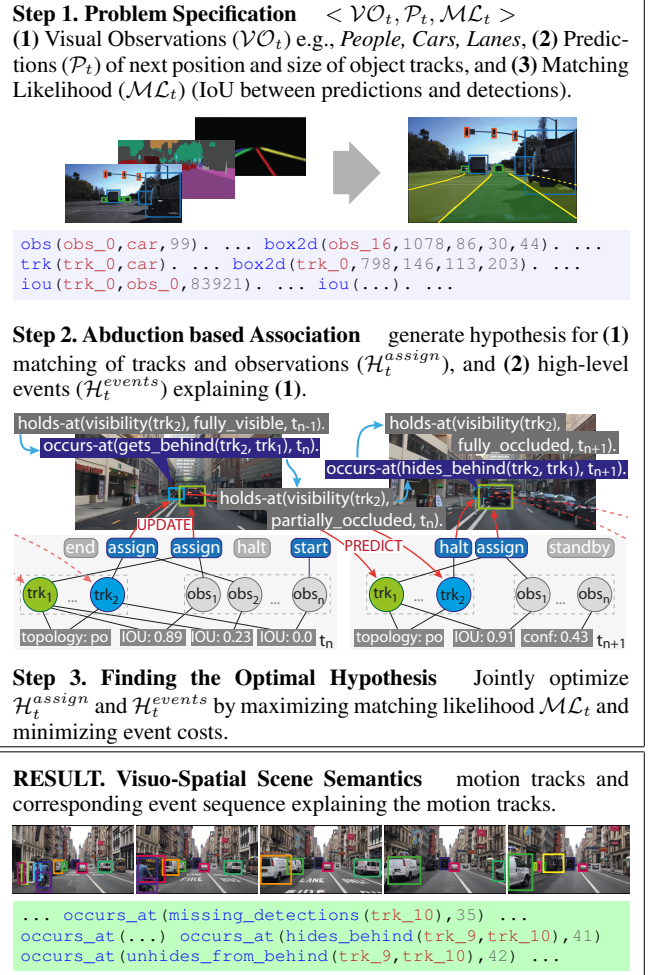


Figure 3: Online Visual Abduction; Comp. Steps for each $t \in T$

etc using YOLOv3 [Redmon and Farhadi, 2018]. Further we generate scene context using *Semantic Segmentation* – segmenting the road, sidewalk, buildings, cars, people, trees, etc. using DeepLabv3+ [Chen *et al.*, 2018], and *Lane Detection* – estimating lane markings, to detect lanes on the road, using SCNN [Pan *et al.*, 2018]. Type and confidence score for each observation is given by $type_{obs_i}$ and $conf_{obs_i}$.

Movement Prediction (\mathcal{P}_t): For each track trk_i changes in *position* and *size* are predicted using kalman filters; this results in an estimate of the spatial entity ε for the next time-point t of each motion track $\mathcal{P}_t = \{\varepsilon_{trk_1}, \dots, \varepsilon_{trk_n}\}$.

Matching Likelihood (\mathcal{ML}_t): For each pair of tracks and observations ε_{trk_i} and ε_{obs_j} , where $\varepsilon_{trk_i} \in \mathcal{P}_t$ and $\varepsilon_{obs_j} \in \mathcal{VO}_t$, we compute the likelihood $\mathcal{ML}_t = \{ml_{trk_i, obs_1}, \dots, ml_{trk_i, obs_j}\}$ that ε_{obs_j} belongs to ε_{trk_i} . The intersection over union (IoU) provides a measure for the amount of overlap between the *spatial entities* ε_{obs_j} and ε_{trk_i} .

Step 2. Abduction based Association Following perception as logical abduction most directly in the sense of Shanahan [2005], we define the task of abducting visual explanations as finding an association (\mathcal{H}_t^{assign}) of observed scene

elements (\mathcal{VO}_t) to the motion tracks of objects (\mathcal{MT}) given by the predictions \mathcal{P}_t , together with a high-level explanation (\mathcal{H}_t^{events}), such that $[\mathcal{H}_t^{assign} \wedge \mathcal{H}_t^{events}]$ is consistent with the background knowledge and the previously abduced event sequence \mathcal{H}^{events} , and entails the perceived scene given by $\langle \mathcal{VO}_t, \mathcal{P}_t, \mathcal{ML}_t \rangle$:

$$\Sigma \wedge \mathcal{H}^{events} \wedge [\mathcal{H}_t^{assign} \wedge \mathcal{H}_t^{events}] \models \mathcal{VO}_t \wedge \mathcal{P}_t \wedge \mathcal{ML}_t$$

where \mathcal{H}_t^{assign} consists of the assignment of detections to object tracks, and \mathcal{H}_t^{events} consists of the high-level events Θ explaining the assignments.

We **associate objects and observations** by finding the best match between observations \mathcal{VO}_t and predicted object tracks \mathcal{P}_t . Towards this we use *choice rules* [Gebser *et al.*, 2014] (i.e., one of the heads of the rule has to be in the stable model) for ε_{obs_j} and ε_{trk_i} , generating all possible assignments in terms of assignment actions: *assign*, *start*, *end*, *halt*, *resume*, *ignore_det*, *ignore_trk*, and then maximising the matching likelihood ml_{trk_i, obs_j} between pairs of spatial entities for matched observations ε_{obs_j} and predicted track region ε_{trk_i} (See Step 3).

```
1{ assign(Trk, Det): det(Det, _, _);
  end(Trk); ignore_trk(Trk); halt(Trk);
  resume(Trk, Det): det(Det, _, _) }1 :- trk(Trk, _).

1{ assign(Trk, Det): trk(Trk, _);
  start(Det); ignore_det(Det);
  resume(Trk, Det): trk(Trk, _). }1 :- det(Det, _, _).
```

For each assignment action we define *integrity constraints*³ that restrict the set of answers generated by the choice rules, e.g., the following constraints are applied to assigning an observation ε_{obs_j} to a track trk_i , applying thresholds on the IoU_{trk_i, obs_j} and the confidence of the observation $conf_{obs_j}$, further we define that the type of the observation has to match the type of the track it is assigned to:

```
:- assign(Trk, Det), not assignment_constraints(Trk, Det).

assignment_constraints(Trk, Det) :-
  trk(Trk, Trk_Type, trk_state(Trk, active),
  det(Det, Det_Type, Conf), Conf > conf_thresh_assign,
  match_type(Trk_Type, Det_Type),
  iou(Trk, Det, IOU), IOU > iou_thresh.
```

Abducible high-level events guide *association* and serve as explanations for object dynamics. For the length of this paper, we restrict to high-level visuo-spatial abducibles pertaining to *object persistence* and *visibility* (Table 1): **(1). Occlusion:** Objects can disappear or reappear as result of occlusion with other objects; **(2). Noise and Missing Observation:** (Missing-)observations can be the result of faulty detections.

Lets take the case of *occlusion*: functional fluent visibility could be denoted *fully_visible*, *partially_occluded* or *fully_occluded*:

```
fluent(visibility(Trk)) :- trk(Trk, _).

possVal(visibility(Trk), fully_visible) :- trk(Trk, _).
possVal(visibility(Trk), partially_visible) :- trk(Trk, _).
possVal(visibility(Trk), not_visible) :- trk(Trk, _).
```

We define the event *hides_behind/2*, stating that an object hides behind another object by defining the conditions

³ Integrity constraints restrict the answer set by eliminating stable models where the body is satisfied.

EVENTS	Description	
hides_behind(Trk ₁ , Trk ₂)	Trk ₁ hides behind Trk ₂ .	
unhides_from_behind(Trk ₁ , Trk ₂)	Trk ₁ unhides from behind Trk ₂ .	
missing_detections(Trk)	Missing detections in Trk.	
FLUENTS	Values	Description
hidden_by(Trk1, Trk2)	{true;false}	Trk1 is hidden by Trk2.
visibility(Trk)	{fully_visible; partially_occluded; fully_occluded}	Visibility state of Trk.

Table 1: Events and Fluents Explaining (Dis)Appearance

that have to hold for the event to possibly occur, and the effects the occurrence of the event has on the properties of the objects, i.e., the value of the visibility fluent changes to *fully_occluded*.

```
event(hides_behind(Trk1, Trk2)) :- trk(Trk1, _), trk(Trk2, _).

causesValue(hides_behind(Trk1, Trk2),
  visibility(Trk1), not_visible, T) :-
  trk(Trk1, _), trk(Trk2, _), time(T).

:- occurs_at(hides_behind(Trk1, Trk2), curr_time),
  trk(Trk1, _), trk(Trk2, _),
  not position(overlapping_top, Trk1, Trk2).
```

For abducing the occurrence of an event we use choice rules that connect the event with assignment actions, e.g., a track getting halted may be explained by the event that the track hides behind another track.

```
1{ occurs_at(hides_behind(Trk, Trk2), curr_time):
  trk(Trk2, _); ... }1 :- halt(Trk).
```

Step 3. Finding the Optimal Hypothesis To ensure an *optimal assignment*, we use ASP based optimization to maximize the matching likelihood between matched pairs of tracks and detections. Towards this, we first define the matching likelihood based on the Intersection over Union (IoU) between the observations and the predicted boxes for each track as described in [Bewley *et al.*, 2016]:

```
assignment_prob(Trk, Det, IOU) :-
  det(Det, _, _), trk(Trk, _), iou(Trk, Det, IOU).
```

We then maximize the matching likelihood for all assignments, using the build in *maximize* statement:

```
#maximize {(Prob)@1, Trk, Det :
  assign(Trk, Det), assignment_prob(Trk, Det, Prob)}.
```

To find the best set of hypotheses with respect to the observations, we *minimize* the occurrence of certain events and association actions, e.g., the following optimization statements minimize starting and ending tracks; the resulting assignment is then used to update the motion tracks accordingly.

```
#minimize {5@2, Trk: end(Trk)}.
#minimize {5@2, Det: start(Det)}.
```

It is important here to note that: **(1).** by jointly abducing the object dynamics and high-level events we can impose constraints on the assignment of detections to tracks, i.e., an assignment is only possible if we can find an explanation supporting the assignment; and **(2).** the likelihood that an event

Situation	Objects	Description
HIDDEN_ENTITY	<i>entity, object</i>	traffic participant hidden by obstacle
REDUCED_VISIBILITY	<i>object</i>	visibility reduced by object in front.
SUDDEN_STOP	<i>vehical</i>	vehicle in front stopping suddenly
BLOCKED_LANE	<i>lane, object</i>	lane of the road is blocked by an object.

Table 2: Safety-Critical Situations

occurs guides the assignments of observations to tracks. Instead of independently tracking objects and interpreting the interactions, this yields to event sequences that are consistent with the abducted object tracks, and noise in the observations is reduced (See evaluation in Sec. 3).

3 Application & Evaluation

We demonstrate applicability towards identifying and interpreting *safety-critical situations* (e.g., Table 2); these encompass those scenarios where interpretation of spacetime dynamics, driving behaviour, environmental characteristics is necessary to anticipate and avoid potential dangers.

Reasoning about Hidden Entities Consider the situation of Fig. 4: a car gets occluded by another car turning left and reappears *in front of* the autonomous vehicle. Using online abduction for abducting high-level interactions of scene objects we can hypothesize that the car got *occluded* and anticipate its reappearance based on the perceived scene dynamics. The following shows data and abducted events.

```

trk(trk_3, car). trk_state(trk_3, active). ...
trk(trk_41, car). trk_state(trk_41, active). ...
box2d(trk_3, 660, 460, 134, 102). ...
box2d(trk_41, 631, 471, 40, 47). ...
... occurs_at(hides_behind(trk_41, trk_3), 179)) ...
    
```

We define a rule stating that a *hidden* object may *unhide* from behind the object it is hidden by and anticipate the time point t based on the object *movement* as follows:

```

anticipate(unhides_from_behind(Trk1, Trk2), T) :-
    time(T), curr_time < T,
    holds_at(hidden_by(Trk1, Trk2), curr_time),
    topology(proper_part, Trk1, Trk2),
    movement(moves_out_of, Trk1, Trk2, T).
    
```

We then interpolate the objects position at time point t to predict where the object may *reappear*.

```

point2d(interpolated_position(Trk, T), PosX, PosY) :-
    time(T), curr_time < T, T1 = T - curr_time,
    box2d(Trk1, X, Y, _, _), trk_mov(Trk1, MovX, MovY),
    PosX = X + MovX * T1, PosY = Y + MovY * T1.
    
```

For the occluded car in our example we get the following prediction for time t and position x, y :

```

anticipate(unhides_from_behind(trk_41, trk_2), 202)
point2d(interpolated_position(trk_41, 202), 738, 495)
    
```

Based on this prediction we can then define a rule that gives a warning if a hidden entity may reappear in front of the vehicle, which could be used by the control mechanism, e.g., to adapt driving and slow down in order to keep safe distance:




Figure 4: Abducting Occlusion to Anticipate Reappearance

```

warning(hidden_entity_in_front(Trk1, T)) :-
    time(T), T - curr_time < anticipation_threshold,
    anticipate(unhides_from_behind(Trk1, _), T),
    position(in_front, interpolated_pos(Trk1, T)).
    
```

Empirical Evaluation For online sensemaking, evaluation focusses on accuracy of abducted motion tracks, real-time performance, and the tradeoff between performance and accuracy. Our evaluation uses the **KITTI object tracking dataset** [Geiger *et al.*, 2012], which is a community established benchmark dataset for autonomous cars: it consists of 21 training and 29 test scenes, and provides accurate track annotations for 8 object classes (e.g., car, pedestrian, van, cyclist). We also evaluate tracking results using the more general cross-domain **Multi-Object Tracking (MOT)** dataset [Milan *et al.*, 2016] established as part of the *MOT Challenge*; it consists of 7 training and 7 test scenes which are highly unconstrained videos filmed with both static and moving cameras. We evaluate on the available groundtruth for training scenes of both KITTI using YOLOv3 detections, and MOT17 using the provided faster RCNN detections.

For **evaluating object tracking** we consider *accuracy* (MOTA) and *precision* (MOTP) of abducted object tracks following the Clear MOT [Bernardin and Stiefelwagen, 2008] evaluation schema. Results (Table 3) show that jointly abducting high-level object interactions together with low-level scene dynamics increases the accuracy of the object tracks, i.e. we consistently observe an improvement of about 5%, from 45.72% to 50.5% for *cars* and 28.71% to 32.57% for *pedestrians* on KITTI, and from 41.4% to 46.2% on MOT.

Online performance and scalability of online abduction is evaluated with respect to its real-time capabilities.⁴ (1). We compare the time & accuracy of online abduction for state of the art (real-time) detection methods: YOLOv3, SSD [Liu *et al.*, 2016], and Faster RCNN [Ren *et al.*, 2015] (Fig. 5). (2). We evaluate scalability of the ASP based abduction on a synthetic dataset with controlled number of tracks and % of overlapping tracks per frame. Results (Fig. 5) show that online abduction can perform with above 30 frames per second for scenes with up to 10 highly overlapping object tracks, and more than 50 tracks with 1fps (for the sake of testing, it is worth noting that even for 100 objects per frame it only takes about an average of 4 secs per frame). Importantly, for realistic scenes such as in the KITTI dataset, abduction runs realtime at 33.9fps using YOLOv3, and 46.7 using SSD with lower accuracy but providing good precision.

Discussion of Empirical Results Results show that integrating high-level abduction and object tracking improves the

⁴Evaluation using a dedicated Intel Core i7-6850K 3.6GHz 6-Core Processor, 64GB RAM, and a NVIDIA Titan V GPU 12GB.

SEQUENCE	Tracking	MOTA	MOTP	ML	MT	FP	FN	ID sw.	Frag.
KITTI tracking – Cars (8008 frames, 636 targets)	without Abduction	45.72 %	76.89 %	19.14 %	23.04 %	785	11182	1097	1440
	with Abduction	50.5 %	74.76 %	20.21 %	23.23 %	1311	10439	165	490
KITTI tracking – Pedestrians (8008 frames, 167 targets)	without Abduction	28.71 %	71.43 %	26.94 %	9.58 %	1261	6119	539	833
	with Abduction	32.57 %	70.68 %	22.15 %	14.37 %	1899	5477	115	444
MOT 2017 (5316 frames, 546 targets)	without Abduction	41.4 %	88.0 %	35.53 %	16.48 %	4877	60164	779	741
	with Abduction	46.2 %	87.9 %	31.32 %	20.7 %	5195	54421	800	904

Table 3: Evaluation of Tracking Performance; accuracy (MOTA), precision (MOTP), mostly tracked (MT) and mostly lost (ML) tracks, false positives (FP), false negatives (FN), identity switches (ID Sw.), and fragmentation (Frag.).

DET.	Recall	MOTA	MOTP	fps_{det}	fps_{abd}
YOLOv3	0.690	50.5 %	74.76 %	45	33.9
SSD	0.599	30.63 %	77.4 %	8	46.7
FRCNN	0.624	37.96 %	72.9 %	5	32.0




Figure 5: Online Performance and Scalability; performance for pre-trained detectors (DET.) on the 'cars' class of KITTI dataset, and processing time relative to the no. of tracks on synthetic dataset.

resulting object tracks and reduce the noise in the visual observations. For the case of online visual sense-making, ASP based abduction provides the required performance: even though the complexity of ASP based abduction increases quickly, with large numbers of tracked objects the framework can track up to 20 objects simultaneously with 30fps and achieve real-time performance on the KITTI benchmark dataset. It is also important to note that the tracking approach in this paper is based on *tracking by detection* using a naive measure, i.e., the IoU (Sec. 2.2; Step 1), to associate observations and tracks, and it is not using any visual information in the prediction or association step. Naturally, this results in a lower accuracy, in particular when used with noisy detections and when tracking fast moving objects in a benchmark dataset such as KITTI. That said, due to the modularity of the implemented framework, extensions with different methods for predicting motion (e.g., using particle filters or optical flow based prediction) are straightforward: i.e., improving tracking is not the aim of our research.

4 Related Work

Answer Set Programming is now widely used as an underlying knowledge representation language and robust methodology for non-monotonic reasoning [Brewka *et al.*, 2011; Gebser *et al.*, 2012]. With ASP as a foundation, and driven by semantics, commonsense and explainability [Davis and Marcus, 2015], this research aims to bridge the gap between high-level formalisms for logical abduction and low-level visual processing by tightly integrating semantic abstractions of space-change with their underlying numerical representations. Within KR, the significance of high-level (abductive)

explanations in a range of contexts is well-established: planning & process recognition [Kautz, 1991], vision & abduction [Shanahan, 2005], probabilistic abduction [Blythe *et al.*, 2011], reasoning about spatio-temporal dynamics [Bhatt and Loke, 2008], reasoning about continuous *spacetime* change [Muller, 1998; Hazarika and Cohn, 2002] etc. Dubba *et al.* [2015] uses abductive reasoning in an inductive-abductive loop within inductive logic programming (ILP). Aditya *et al.* [2015] formalise general rules for image interpretation with ASP. Similarly motivated to this research is [Suchan *et al.*, 2018], which uses a two-step approach (with one huge *problem specification*), first tracking and then explaining (and fixing) tracking errors; such an approach is not runtime / real-time capable. In computer vision research there has recently been an interest to synergise with cognitively motivated methods; in particular, e.g., for perceptual grounding and inference [Yu *et al.*, 2015], and combining video analysis with textual information for understanding events and answering queries about video data [Tu *et al.*, 2014].

5 Conclusion & Outlook

We develop a novel abduction-driven *online* (i.e., realtime, incremental) visual sense-making framework: general, systematically formalised, modular and fully implemented. Integrating robust state-of-the-art methods in *knowledge representation* and *computer vision*, the framework has been evaluated and demonstrated with established benchmarks. We highlight application prospects of semantic vision for autonomous driving, a domain of emerging and long-term significance. Specialised commonsense theories (e.g., about multi-sensory integration and multi-agent belief merging, contextual knowledge) may be incorporated based on requirements. Our ongoing focus is to develop a novel dataset emphasising semantics and (commonsense) explainability; this is driven by mixed-methods research –AI, Psychology, HCI– for the study of driving behaviour in low-speed, complex urban environments with unstructured traffic. Such interdisciplinary studies are needed to better appreciate the complexity and spectrum of varied human-centred challenges in autonomous driving, and demonstrate the significance of integrated vision and semantics solutions in those contexts.

Acknowledgements

Partial funding by the German Research Foundation (DFG) via the CRC 1320 EASE – Everyday Activity Science and Engineering” (www.ease-crc.org), Project P3: *Spatial Reasoning in Everyday Activity* is acknowledged.

References

- [Aditya *et al.*, 2015] Somak Aditya, Yezhou Yang, Chitta Baral, Cornelia Fermuller, and Yiannis Aloimonos. Visual common-sense for scene understanding using perception, semantic parsing and reasoning. In *2015 AAAI Spring Symposium Series*, 2015.
- [Allen, 1983] James F. Allen. Maintaining knowledge about temporal intervals. *Commun. ACM*, 26(11):832–843, 1983.
- [Balbiani *et al.*, 1999] Philippe Balbiani, Jean-François Condotta, and Luis Fariñas del Cerro. A new tractable subclass of the rectangle algebra. In Thomas Dean, editor, *IJCAI 1999, Sweden*, pages 442–447. Morgan Kaufmann, 1999.
- [Bernardin and Stiefelwagen, 2008] Keni Bernardin and Rainer Stiefelwagen. Evaluating multiple object tracking performance: The clear mot metrics. *EURASIP Journal on Image and Video Processing*, 2008(1):246309, May 2008.
- [Bewley *et al.*, 2016] Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Upcroft. Simple online and realtime tracking. In *2016 IEEE International Conference on Image Processing (ICIP)*, pages 3464–3468, 2016.
- [Bhatt and Loke, 2008] Mehul Bhatt and Seng W. Loke. Modelling dynamic spatial systems in the situation calculus. *Spatial Cognition & Computation*, 8(1-2):86–130, 2008.
- [Bhatt *et al.*, 2013] Mehul Bhatt, Carl Schultz, and Christian Freksa. The ‘Space’ in Spatial Assistance Systems: Conception, Formalisation and Computation. In: *Representing space in cognition: Interrelations of behavior, language, and formal models. Series: Explorations in Language and Space*. 978-0-19-967991-1, Oxford University Press, 2013.
- [Blythe *et al.*, 2011] James Blythe, Jerry R. Hobbs, Pedro Domingos, Rohit J. Kate, and Raymond J. Mooney. Implementing weighted abduction in markov logic. In *Proc. of 9th Intl. Conference on Computational Semantics, IWCS ’11, USA, 2011*. ACL.
- [BMVI, 2018] BMVI. Report by the ethics commission on automated and connected driving., BMVI: Federal ministry of transport and digital infrastructure, germany, 2018.
- [Brewka *et al.*, 2011] Gerhard Brewka, Thomas Eiter, and Miroslaw Truszczyński. Answer set programming at a glance. *Commun. ACM*, 54(12):92–103, December 2011.
- [Chen *et al.*, 2018] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. *arXiv:1802.02611*, 2018.
- [Davis and Marcus, 2015] Ernest Davis and Gary Marcus. Commonsense reasoning and commonsense knowledge in artificial intelligence. *Commun. ACM*, 58(9):92–103, 2015.
- [Dubba *et al.*, 2015] Krishna Sandeep Reddy Dubba, Anthony G. Cohn, David C. Hogg, Mehul Bhatt, and Frank Dylla. Learning relational event models from video. *J. Artif. Intell. Res. (JAIR)*, 53:41–90, 2015.
- [Gebser *et al.*, 2012] Martin Gebser, Roland Kaminski, Benjamin Kaufmann, and Torsten Schaub. *Answer Set Solving in Practice*. Morgan & Claypool Publishers, 2012.
- [Gebser *et al.*, 2014] Martin Gebser, Roland Kaminski, Benjamin Kaufmann, and Torsten Schaub. Clingo = ASP + control: Preliminary report. *CoRR*, abs/1405.3694, 2014.
- [Geiger *et al.*, 2012] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [Hazarika and Cohn, 2002] Shyamanta M Hazarika and Anthony G Cohn. Abducing qualitative spatio-temporal histories from partial observations. In *KR*, pages 14–25, 2002.
- [Kautz, 1991] Henry A. Kautz. Reasoning about plans. chapter A Formal Theory of Plan Recognition and Its Implementation, pages 69–124. Morgan Kaufmann Publishers Inc., USA, 1991.
- [Liu *et al.*, 2016] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott E. Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: single shot multibox detector. In *ECCV (1)*, volume 9905 of *LNCS*, pages 21–37. Springer, 2016.
- [Ma *et al.*, 2014] Jiefei Ma, Rob Miller, Leora Morgenstern, and Theodore Patkos. An epistemic event calculus for asp-based reasoning about knowledge of the past, present and future. In *LPAR: 19th Intl. Conf. on Logic for Programming, Artificial Intelligence and Reasoning*, volume 26 of *EPiC Series in Computing*, pages 75–87. EasyChair, 2014.
- [Mani and Pustejovsky, 2012] Inderjeet Mani and James Pustejovsky. *Interpreting Motion - Grounded Representations for Spatial Language*, volume 5 of *Explorations in language and space*. Oxford University Press, 2012.
- [Milan *et al.*, 2016] Anton Milan, Laura Leal-Taixé, Ian D. Reid, Stefan Roth, and Konrad Schindler. MOT16: A benchmark for multi-object tracking. *CoRR*, abs/1603.00831, 2016.
- [Miller *et al.*, 2013] Rob Miller, Leora Morgenstern, and Theodore Patkos. Reasoning about knowledge and action in an epistemic event calculus. In *COMMONSENSE 2013*, 2013.
- [Muller, 1998] Philippe Muller. A qualitative theory of motion based on spatio-temporal primitives. In Anthony G. Cohn et al., editor, *KR 1998, Italy*. Morgan Kaufmann, 1998.
- [Pan *et al.*, 2018] Xingang Pan, Jianping Shi, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Spatial as deep: Spatial CNN for traffic scene understanding. In Sheila A. McIlraith and Kilian Q. Weinberger, editors, *AAAI 2018*. AAAI Press, 2018.
- [Redmon and Farhadi, 2018] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *CoRR*, abs/1804.02767, 2018.
- [Ren *et al.*, 2015] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. In *Annual Conference on Neural Information Processing Systems 2015, Canada, 2015*.
- [Shanahan, 2005] Murray Shanahan. Perception as abduction: Turning sensor data into meaningful representation. *Cognitive Science*, 29(1):103–134, 2005.
- [Suchan and Bhatt, 2016] Jakob Suchan and Mehul Bhatt. Semantic question-answering with video and eye-tracking data: AI foundations for human visual perception driven cognitive film studies. In S. Kambhampati, editor, *IJCAI 2016, New York, USA*, pages 2633–2639. IJCAI/AAAI Press, 2016.
- [Suchan *et al.*, 2018] Jakob Suchan, Mehul Bhatt, Przemyslaw Andrzej Walega, and Carl P. L. Schultz. Visual explanation by high-level abduction: On answer-set programming driven reasoning about moving objects. In: *AAAI 2018*. AAAI Press, 2018.
- [Tu *et al.*, 2014] Kewei Tu, Meng Meng, Mun Wai Lee, Tae Eun Choe, and Song Chun Zhu. Joint video and text parsing for understanding events and answering queries. *IEEE MultiMedia*, 2014.
- [Yu *et al.*, 2015] Haonan Yu, N. Siddharth, Andrei Barbu, and Jeffrey Mark Siskind. A Compositional Framework for Grounding Language Inference, Generation, and Acquisition in Video. *J. Artif. Intell. Res. (JAIR)*, 52:601–713, 2015.