

Plants Don't Walk on the Street: Common-Sense Reasoning for Reliable Semantic Segmentation

Linara Adilova^{1,*}, Elena Schulz^{1,*}, Maram Akila¹, Sebastian Houben¹,
Jan David Schneider², Fabian Hüger², Tim Wirtz¹
¹ Fraunhofer IAIS, ² Volkswagen AG

* equal contribution

{first.last}@iais.fraunhofer.de

fabian.hueger@volkswagen.de

jan.david.schneider@volkswagen.de

Abstract

Data-driven sensor interpretation in autonomous driving can lead to highly implausible predictions as can most of the time be verified with common-sense knowledge. However, learning common knowledge only from data is hard and approaches for knowledge integration are an active research area. We propose to use a partly human-designed, partly learned set of rules to describe relations between objects of a traffic scene on a high level of abstraction. In doing so, we improve and robustify existing deep neural networks consuming low-level sensor information. We present an initial study adapting the well-established Probabilistic Soft Logic (PSL) framework to validate and improve on the problem of semantic segmentation. We describe in detail how we integrate common knowledge into the segmentation pipeline using PSL and verify our approach in a set of experiments demonstrating the increase in robustness against several severe image distortions applied to the A2D2¹ autonomous driving data set.

1. Introduction

Deep neural networks are an integral part of autonomous driving as they achieve unmatched performance in a number of tasks, in particular image-based environment perception. Still, they are prone to random or targeted disturbances, which renders robustness a crucial requirement for safety-critical applications. These disturbances might be of a natural source – like overexposure, fog, snow – or of an artificial and purposefully malicious source – like adversarial attacks. At the same time, wrong predictions in such tasks are most of the time blatantly obvious for hu-

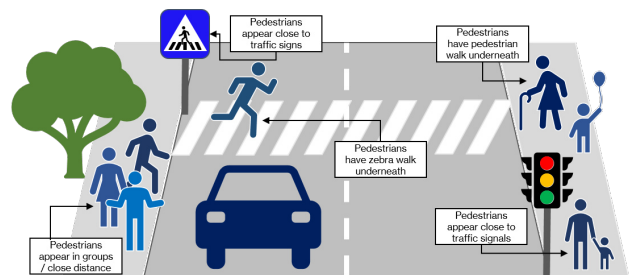


Figure 1: Most humans approach the task of understanding and reasoning about traffic scenes using relations between objects. Given that all of them have been identified, common knowledge — describing both behaviour and features, e.g. relative position — can straightforwardly be delineated in terms of simple rules. Some of the rules that we utilize in this paper are exemplified in the above figure.

man viewers that use multiple intuitive rules based on their knowledge and experience to reveal what is happening in a given scene. Neural approaches, in their current form, aim to implicitly learn these rules from data, which makes it necessary to acquire and annotate large amounts of inputs and furthermore introduces the danger of adopting spurious correlations. The integration of scene knowledge into the training or inference of neural networks seems a viable approach, however, this knowledge ranges from considerations of scene geometry, common knowledge, and rules of human behavior, rendering it a difficult problem [17]. There are not only a number of competing approaches to perform the integration of knowledge itself but also multiple ways of representing the knowledge.

In this paper, we study a way to represent knowledge in form of relations delineated as logic rules between identified objects in a traffic scene (see Fig. 1) and examine how

¹<https://www.a2d2.audi/a2d2/en.html>

these relations can be used to robustify data-driven environment perception. To this end, we adapt the Probabilistic Soft Logic (PSL) framework [10] to define and learn relations between objects in the surrounding scene. We demonstrate the results of initial experiments aimed at quantifying its effect on the robustness of a recent small-scale semantic segmentation neural network.

2. Related Work

The importance of high-level knowledge injection for image-related tasks is widely accepted [1]. Utilizing background and common-sense knowledge about objects or regions, which goes beyond data annotations, can benefit many computer vision tasks, such as object recognition or semantic segmentation. Mostly due to its use in robotics and autonomous driving vehicles, semantic segmentation, as one form of low-level scene representation, is currently one of the actively developing fields in image recognition.

In their survey Siam et al. [15] show the development of semantic segmentation from feature-based approaches to the end-to-end deep learning networks. Among other techniques, they point out and discuss an increase in accuracy achieved by including structural knowledge, which is abundant in the particular task of autonomous driving. To facilitate knowledge inclusion, according to Siam et al. [15], many researchers apply conditional random fields (CRF) additionally to deep convolutional networks for classification to improve the accuracy of the predictions. There are also approaches that are based only on an ontology or logical reasoning for semantic segmentation. For example, Zand et al. [18] propose to deploy Dirichlet mixture models based on super-pixel segmentation using knowledge from a fixed, user-defined ontology after which prediction is performed using conditional random fields. Analogously, Leon-Garza et al. [12] use fuzzy logic to perform segmentation and Dasiopoulou et al. [4] use a high-level semantic ontology for object detection in videos.

In our work, we propose to employ the PSL framework as an expressive and interpretable way to reason about the environment of the autonomous vehicle. To the best of our knowledge it is the first attempt to combine deep-learning-based semantic segmentation with probabilistic logical reasoning about the environment. Nevertheless, the idea of employing knowledge via logic rules in autonomous driving tasks is not new. In fact, Kanaujia et al. [7], Kardaş et al. [8], Kembhavi et al. [9] use Markov logic networks for recognizing complex events in video data.

Related work in the same area of application as ours is by Souza and Santos [16], in which the authors use logic reasoning to support a lane detection system with high-level knowledge. They employ the output of their segmentation method for further reasoning and controlling the movement of the vehicle (or alerting the driver). Similar to the PSL

framework used in our work, these authors adopt Markov logic networks to employ high-level knowledge in the form of logical hard and soft rules. However, the object-based rules are only used to classify events after recognizing objects (*e.g.* person, car) within image frames. In contrast, we employ soft object-based rules to robustify the prediction results of a semantic segmentation neural network for the objects in the image frames itself. Thus, our goal is to construct rules that analogously would support semantic segmentation in cases when the network itself is compromised (*e.g.* by an image distortion) in order to prevent completely missing critical objects such as pedestrians, traffic signs or traffic signals.

3. Approach

The goal of our research is to find ways in which high-level domain knowledge can improve and robustify the predictions of deep learning models in the context of autonomous driving. While multiple lines of research discussed in Sec. 2 make use of domain knowledge exclusively for building a model, it is our aim to combine the data-driven approach and high-level knowledge during inference.

For various traffic scenes, humans can typically use common and empirical knowledge about behaviour and features of the objects to infer their type and in-scene relations, even if generalizing to unseen scenarios. Considering different features and interrelations of objects, this human common knowledge can partially be encoded as a set of simple rules, some of which are showcased in Fig. 1. The way a human would reason about such a traffic scene cannot be based on hard Boolean logic, though, in which rules hold either always or never. For example, an object having a zebra walk underneath and standing close to a traffic sign has a high probability to be a human, as depicted in Fig. 1. But there is a chance that this could also be something else, *e.g.* a dog or some obstacle in a construction site. Therefore, to integrate such high-level human knowledge in the same “uncertain” manner into a neural network model, we need to define logical rules with relaxed truth values. Structuring the rules upon the identification and classification of objects aligns with the task of semantic segmentation that lends itself well to finding image regions with a common semantic, which we treat as object proposals. We selected PSL as a framework for knowledge representation for two reasons: First, it allows us to flexibly use soft truth values that are suitable to define both strict and loose rules, and, second, it can be fine-tuned on data, hence, distilling high-level knowledge directly via the given annotations.

We first introduce PSL in Sec. 3.1 and briefly describe the input requirements for a PSL program. Based on this, Sec. 3.2 elaborates on how we define object entities and rules in our domain of interest followed by a description

of the approach to combine the PSL rules with semantic segmentation.

3.1. Probabilistic Soft Logic

Probabilistic Soft Logic (PSL) [3, 10] is a programming framework for probabilistic reasoning in domains with relational structure. Typical applications of PSL are collective classification (of interlinked objects in structured data), entity resolution (finding a mapping from multiple interrelated references to the true set of underlying entities) and link prediction (predicting whether a relationship exists between two entities). PSL makes use of Hinge-Loss Markov Random Fields (HL-MRF), a graphical model, and applies these to structured machine learning problems. Graphical models define a type of loss function that represents a probability score for each possible value configuration respecting the graph-encoded dependencies among the variables. For HL-MRFs [2, 3] the continuous loss functions are defined using the hinge loss, which allows for efficient convex optimization techniques during inference.

The HL-MRFs are defined using declarative rules, whose importance can be indicated by optional non-negative *weights*. These rules are formulated using first order logic. In contrast to the binary truth values *true* and *false* of Boolean algebra, PSL uses continuous soft truth values in the interval $[0, 1]$. For a detailed explanation of the elements of a PSL rule, see the example depicted in Fig. 2.

Conceptually, we require that the “satisfaction” of rules can be measured in a soft-Boolean way if all “variables” have a defined value. In this case, we can aggregate the set of all rules to a global score using individual weights for each rule. Given a set of “ground truth” observations, the weights can be fine-tuned such that the importance of each rule agrees with the ground truth observations. The resulting fixed set of weights then allows us to perform inference over incomplete observations by choosing the most probable value for unknown variables.

In detail, this means each atom is mapped to a soft truth value – before inference for all known variables, and during inference for all unknown variables. To compute the continuous truth value of the whole rule, values on atom-level are combined using the *Lukasiewicz t-norm* [11] as a relaxation for the logical conjunction, disjunction, and negation. Each rule’s *satisfaction* is calculated indicating to what degree that rule is fulfilled given the current configuration of values. A rule is satisfied if the head is assigned at least the same truth value as the body, and a fully satisfied rule has a value of 1.

All possible variable configurations for the rules are combined to a probability distribution, in which those that satisfy more ground rules are more probable. If a weight is assigned to a rule, this indicates the importance of this rule to be satisfied in relation to the other rules and their corre-

sponding weights. Rules that are unweighted induce hard constraints which require the rule to be met at all times. During inference, the most probable value configuration for unknown predicates is found given evidence (via known predicates). Given ground truth, the weights of rules can also be learned with maximum-likelihood estimation. Thus, the importance of each rule in the rule set can be learned from the data itself.

Apart from these rules, a PSL program consists of a database that defines the elements of the domain of interest as well as their relations in the form of predicates. The database is built up from known predicates (*i.e.* atoms that are completely observed) and unknown predicates (*i.e.* atoms that are unobserved and must be inferred) that are part of one of the three following sets: (i) Necessary for all known predicates is the set of *observations*. These are variable combinations with an associated soft truth value in the form of actual extracted and known data for each. (ii) Necessary for all unknown predicates is the set of *targets* that is a list of all atoms (combinations of variables) that should be inferred, therefore without any associated truth values. (iii) For weight learning a set of *truths* with actual known data (likewise observations) to be used as ground truth for all unknown predicates is necessary. Inference in PSL is done based on the first two sets, while evaluation of the inferred predicates can be done with the test set of the *truths*.

There are two types of rules that are supported in PSL: logical and arithmetic rules. We use logical rules to encode dependencies between atoms as for example depicted in the rule in Fig. 2. Additionally, we need to imply an arithmetic constraint to restrict the substitution for variables further. In our setting, we want to reason about the class of an object in the scene that can only be one of three classes, constraining the corresponding truth values to sum up to 1.

3.2. Deriving Rules from Domain Knowledge

The overall goal of employing PSL to camera-based perception in autonomous driving is to identify parts of common knowledge about the traffic scene that can be readily extracted and represent it in a form that facilitates integration into the model performing inference. Based on the input requirements for a PSL program as defined above, we first need a database, *i.e.* objects that we want to reason about as well as their relationships encoded within chosen predicates. Then, we define a set of rules that carries additional relational knowledge about the domain of interest that is not yet covered by the annotated data and thus can be deployed to support and validate the perception task. When defining an initial set of rules, one therefore has to carefully choose the objects they link up: it should be simple to both describe their behavior and to extract them from raw data.

Here, we choose pedestrians, traffic signs and traffic signals as our three types of objects of interest, for which we




Figure 2: PSL rules have a conjuncted body and a single literal head. *Atoms* (e.g. `IntersectionUnder(01, 'r')`) are the basic elements that are used for reasoning. An atom is a combination of *predicates* (here: `IntersectionUnder\2`) and a sequence of variables and constants equal to the predicate’s arity. A *constant* indicates a concrete element (e.g. “ped” is the concrete object class pedestrian) and a *variable* denotes a placeholder later to be substituted by concrete elements. Here, the predicate `IntersectionUnder` has an arity of two, thus, the atom consists of the variables `01` indicating an object plus the constant “r” indicating the underground type is the road. The identifier `01` will be substituted for all available objects during inference while the constant “r” will stay the same. The `01` identifier is used in three atoms indicating that during inference, each object-1-variable will be substituted by the same constant (same holds for object-2-variable `02`, respectively). The exponent `^2` at the end of the rule indicates that the quadratic loss is chosen. This example rule can be interpreted as “if an object 1 is placed on a sidewalk and in close distance to another object 2, which is a traffic signal, there is a probability that this object 1 is a pedestrian”.

want to infer the class. To identify these objects in a logical setting, one would need to fall back to low-level sensor interpretation. Typically, autonomous vehicles have diverse sensors from which objects, their attributes (e.g. colors, shapes) and their surroundings (e.g. area underneath, background) and their relationships (e.g. distance to other objects) can be derived.

For our case study, however, we solely focus on the aspect of how external knowledge can be meaningfully integrated into an existing prediction workflow. We therefore purposefully deploy an oracle to extract the object entities thereby neglecting additional challenges resulting from, e.g., unreliable sensor readings and object extraction. Our solution is to employ the connected components on ground truth segmentation masks using additional smoothing to avoid fractured objects due to occlusion. Additionally, we emulate a depth sensor by propagating the RGB inputs to a publicly available² version of `MonoDepth2` [6]. This “oracle sensor” information already allows for creating a meaningful set of objects as basis for the rules.

Subsequently, the rules describing the objects have to be defined. To this end, we develop an initial set of rules, for which we refine the weights later on. As we want to reason about the objects of our three selected classes, we define our target predicate `isType(Obj, Class)`. This is used to later infer a soft truth value for each object and the possible three classes, resulting in one class being more likely than the others. We then base the logical reasoning for an object class on the semantic features, on the surroundings of the objects, and on their relations. Surroundings of the objects are described with the observations (known predicates) `intersectionUnder(Obj, Class)` and

`intersectionBehind(Obj, Class)`. For example, to infer the class of an object to be “traffic sign” `intersectionUnder` should have a high likelihood for the classes “pedestrian walk” or “street”. We integrate color as a semantic feature as it is especially important for traffic signs and traffic signals with the predicate `hasColor(Obj, Color)`. Rules we defined using the color predicate are for example inspired by the officially specified color combinations of common German traffic signs. The relations between objects are specified as a distance value based on depth readings with the known predicate `distance(Obj1, Obj2)`. The defined rules for our setup can be found in appendix 6.1.

As a final step, the knowledge gained from the set of data via probabilistic rules has to be combined with the predictions from the neural network, in this case the semantic segmentation masks. For our conceptual demonstration we follow a simplified approach integrating scores obtained from the network as prior assumptions within the rule based framework.

Overall, the practical approach can be described as follows (see Fig. 3):

1. Identify objects to reason about and extract them from sensor data (oracle using ground truth semantic segmentation);
2. Construct the set of rules using expert and common knowledge;
3. Prepare the required information from lower level sensors for the rules inference (depth, colors);
4. Learn the weights of the rules using a training set of recorded camera frames;

²<https://github.com/nianticlabs/monodepth2>

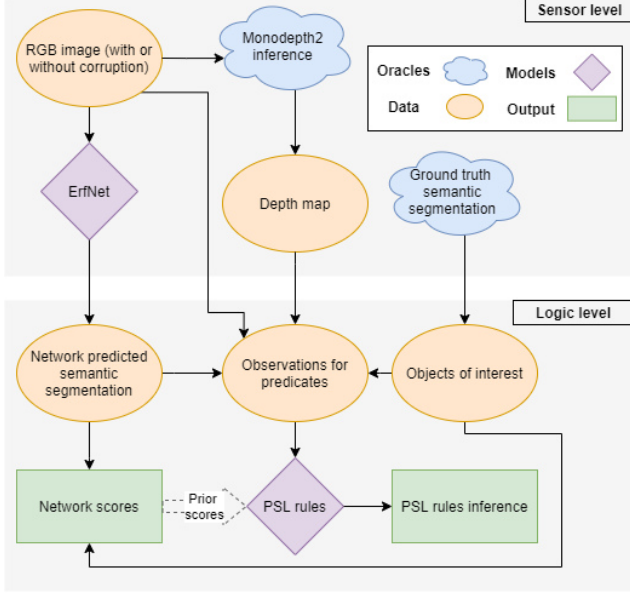


Figure 3: Overall flow of information: Oracles are serving as a source of information for objects and depth maps and also observations for weights learning; predictions produced by the neural network are combined with PSL rules via priors.

5. Perform inference with the neural network and with the rules on the test set of frames for validation, possibly using network outputs as prior.

4. Evaluation

The model that we selected for semantic segmentation is ERFNet [14] whose lightweight encoder-decoder architecture and limited performance make it straightforward to observe the benefit of knowledge injection via PSL. We trained ERFNet on the A2D2 public data set [5], which features a large set of environmental scenes for cars on the road. The data set was separated into independent subsets (by the sequences recorded) for training and testing. The labels in the data set were simplified to 13 classes via grouping. The network was trained using the Adam optimizer with decaying learning rate for 200 epochs. The network achieved an mIoU of 0.59 after training—in particular 0.26 for ‘Pedestrian’ class and 0.63 for ‘Traffic sign’ and ‘Traffic signal’.

Rules selected for the evaluation first have to be ranked with weights. In our evaluation we obtained the weights using the learning procedure, in which the weights are assigned to the rules based on the training set with ground truth annotations. The observations for weight learning are extracted from the training images (predicate `hasColor`), their depth information (predicate `distance`) and the ground truth seman-

tic segmentation (predicates `intersectionBehind` and `intersectionUnder`) using in each case the object entities from the oracle. The full list of rules we defined along with the learned weights can be found in appendix 6.1. This step of weight learning serves to validate our initially proposed rules discarding those which turned out to be inadequate given the data. In these cases, the weights assigned were in the order of 10^{-5} . One such rule was for example increasing the probability of an object to be a pedestrian if it is close to an object that is a traffic sign.

An independent test set of images is used both for evaluation of the network prediction and rule inference. The connected components selected as objects for soft rule-based reasoning are compared against the network predictions for the same region. These regions are produced using the oracle. We refer to the ratio of correctly classified pixels as the network score for this object. Note that for weight learning we used information from the ground truth while for inference network predictions are employed.


To evaluate the robustness of the rule predictions we employed image corruptions as described by Michaelis et al. [13]. As the most natural we selected four, namely brightness, fog, frost and snow. Each of them can be applied with five levels of severity, of which 1 is the weakest (see Fig. 4 for examples). The evaluation then follows the same route – the network infers predictions on the corrupted images and the scores for the objects are computed. Likewise, rules are inferred on corrupted images and their respective predicted semantic segmentation on them. The oracle as further source of information is left unaffected by the applied corruptions. The results of the predictions are depicted in




Figure 4: Example of corruptions: original image (top left), brightness, frost and fog corruption (clockwise).

Fig. 5. We display the scores for the pedestrian objects, either denoting the correct or incorrect class, in the case of the original images and images corrupted with brightness. The histograms for the other two object classes is shown in

the appendix 6.2. One can observe that the scores produced by the rules are nearly not affected—judging by the distribution of histogram peaks and the average score displayed in the legend—while network scores are very sensitive to the corruption of the data.



(a) Scores distribution for the objects “pedestrians” on the original images.



(b) Scores distribution for the objects “pedestrians” on the corrupted (brightness, severity 5) images.

Figure 5: The scores of the network compared to the scores of the rules on the test dataset for pedestrian objects. It should be noted that the distribution of objects is unbalanced in the test set: the number of objects “pedestrian” is the smallest, followed by “traffic signal”. The most frequent class of objects is “traffic sign”. The network was not trained to distinguish between traffic signs or signals and we thus only report the (joined) result as “tr_sign” for the wrong classes.

The final evaluation step is to combine network predictions and rules inference. We used the scores produced by the network as a prior for the `isType` predicate in the PSL rules. Fig. 6 reflects the change in the scores assigned to the objects under increasing severity of corruption. The scores for the objects of class “traffic signal” are presented here, the other classes and corruptions can be seen in the appendix 6.3. One can observe that the severity of the corruption affects the scores produced by a network while the inference with the rules deteriorates only slightly. Finally, the robustification of the network scores is also stable and shows better results than network scores on their own.

5. Discussion and Conclusion

We propose a way to reason about high level scene properties, *i.e.* the type or class of objects, in terms of Probabilis-




Figure 6: The development of average scores for the class “traffic signal” with increasing severity of the “frost” corruption.

tic Soft Logic (PSL). Taking advantage of human knowledge, we draw information from various domains, ranging from sensor input such as color to scene properties such as the relations and distances among object. This approach enables us to include additional information from data-driven model predictions, *e.g.* as prior assumption, and improve their prediction reliability via reasoning. Our initial evaluation in a simplified setup features a scenario from autonomous driving and indicates that the approach is viable. Even in tests with highly adverse input augmentations it is possible to robustify the predictions of the employed semantic segmentation network.

Within the PSL framework reasoning is based on rules, *i.e.* human understandable concepts regarding the interpretation and meaning of object relations or properties. Besides the possibility to guide the rule creation by human expert knowledge this offers strong interpretability, the set of rules could, *e.g.*, be audited with respect to consistency. The relevance of each rule can be adjusted individually, where the necessary weights are either deliberately chosen or learned from a training data set. The latter case, as demonstrated in our evaluation, can serve to validate the applicability of given rules. This places PSL in-between human designed and data driven approaches, making it an especially suited candidate to incorporate common knowledge into fundamentally data driven applications.

This work showcases the potential benefits such common knowledge can have for the reliability of a machine learning framework. At this early stage we, however, neglected the cumbersome step to obtain the entities the PSL framework reasons about. In practice this information has to be aggregated from various, most likely not fully reliable, sources, for instance stemming from additional sensors (*e.g.* LIDAR) or alternative processing approaches. While a probabilistic framework seems appropriate to deal with further uncertain input we restrict this demonstration to only

a single ML model providing the semantic segmentation mask. The entities and their further attributes are instead given by an oracle using the ground truth annotation. This allows us to test the approach without the necessity to instantiate a full framework. Another important step is the involvement of human experts for rules formulation. The presented rules are formulated with the common knowledge of the driving situations and environment, but setup specific rules, as well as the rules that are tuning a particular aspect of recognition need multiple iterations over the rules construction process.

With respect to the incorporation of machine learning models we currently only use the provided prediction. Given the probabilistic nature of both these outputs and the PSL framework the inclusion of prediction uncertainty in terms of, *e.g.*, calibrated softmax scores could improve the reliability further. Adding further models as discussed above would additionally require a larger set of rules. While the used set is sufficient for demonstration purposes enlarging it to cover more attributes seems to be a straightforward way to extend our work. Specific follow up questions could address a more systematic approach to rule building, which investigates the interplay between the set coverage, reflected by the amount and variety of rules, and its achieved performance, reflected in learned weights and task performance, more closely. For the case of autonomous driving the incorporation of temporal consistency, requiring, *e.g.*, that a pedestrian on two consecutive frames should not unexpectedly vanish, could add new aspects to the discussion, possibly including reasoning on the trajectories and intentions of road users.

Especially for domains such as automated driving we believe that the inclusion of human expertise into predictions is a valuable addition to safety and reliability. Often humans are capable to infer relevant information from context. For instance a sign indicating the potential presence of children could, by formulation of an appropriate rule, increase the likelihood that a discovered entity is inferred as child. This inclusion of external knowledge might additionally help counterbalance the rareness of safety relevant situations, which would otherwise make it difficult to learn relevant correlations purely from data. Deliberately fixing weights or the inclusion of specific rules as above might therefore increase the safety of vulnerable road users (VRUs) without affecting the training and optimization of the underlying machine learning models.

Acknowledgements The authors want to thank for support and useful discussions Birgit Kirsch, Julia Rosenzweig, and Dorina Weichert.

The work of L.A. was supported by the Fraunhofer Center for Machine Learning within the Fraunhofer Cluster for Cognitive Internet Technologies. The research of S.H. was

partly funded by the Federal Ministry of Education and Research of Germany as part of the competence center for machine learning ML2R (01IS18038B).

References

- [1] S. Aditya, Y. Yang, and C. Baral. Integrating knowledge and reasoning in image understanding. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 2019.
- [2] S. H. Bach, B. Huang, B. London, and L. Getoor. Hinge-loss markov random fields: Convex inference for structured prediction. *CoRR*, abs/1309.6813, 2013. URL <http://arxiv.org/abs/1309.6813>.
- [3] S. H. Bach, M. Broecheler, B. Huang, and L. Getoor. Hinge-loss markov random fields and probabilistic soft logic. *J. Mach. Learn. Res.*, 18(1):3846–3912, Jan. 2017. ISSN 1532-4435.
- [4] S. Dasiopoulou, V. Mezaris, I. Kompatsiaris, V.-K. Papastathis, and M. G. Strintzis. Knowledge-assisted semantic video object detection. *IEEE Transactions on Circuits and Systems for Video Technology*, 15(10):1210–1224, 2005.
- [5] J. Geyer, Y. Kassahun, M. Mahmudi, X. Ricou, R. Durgesh, A. S. Chung, L. Hauswald, V. H. Pham, M. Mühlegg, S. Dorn, T. Fernandez, M. Jänicke, S. Mirashi, C. Savani, M. Sturm, O. Vorobiov, M. Oelker, S. Garreis, and P. Schuberth. A2d2: Audi autonomous driving dataset, 2020.
- [6] C. Godard, O. Mac Aodha, M. Firman, and G. J. Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3828–3838, 2019.
- [7] A. Kanaujia, T. E. Choe, and H. Deng. Complex events recognition under uncertainty in a sensor network. *CoRR*, abs/1411.0085, 2014. URL <http://arxiv.org/abs/1411.0085>.
- [8] K. Kardaş, İ. Ulusoy, and N. K. Çiçekli. Learning complex event models using markov logic networks. In *2013 IEEE International Conference on Multimedia and Expo Workshops (ICMEW)*, pages 1–6. IEEE, 2013.
- [9] A. Kembhavi, T. Yeh, and L. S. Davis. Why did the person cross the road (there)? scene understanding using probabilistic logic models and common sense reasoning. In *European Conference on Computer Vision*, pages 693–706. Springer, 2010.
- [10] A. Kimmig, S. Bach, M. Broecheler, B. Huang, and L. Getoor. A short introduction to probabilistic soft logic. In *Proceedings of the NIPS Workshop on Probabilistic Programming: Foundations and Applications*, pages 1–4, 2012.
- [11] G. Klir and B. Yuan. *Fuzzy Sets and Fuzzy Logic: Theory and Applications*. Prentice Hall PTR, 1995. ISBN 9780131011717.

- [12] H. Leon-Garza, H. Hagra, A. Peña-Rios, A. Conway, and G. Owusu. A big bang-big crunch type-2 fuzzy logic system for explainable semantic segmentation of trees in satellite images using hsv color space. In *2020 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*, pages 1–7. IEEE, 2020.
- [13] C. Michaelis, B. Mitkus, R. Geirhos, E. Rusak, O. Bringmann, A. S. Ecker, M. Bethge, and W. Brendel. Benchmarking robustness in object detection: Autonomous driving when winter is coming. *arXiv preprint arXiv:1907.07484*, 2019.
- [14] E. Romera, J. M. Alvarez, L. M. Bergasa, and R. Arroyo. Erfnet: Efficient residual factorized convnet for real-time semantic segmentation. *IEEE Transactions on Intelligent Transportation Systems*, 19(1):263–272, 2017.
- [15] M. Siam, S. Elkerdawy, M. Jagersand, and S. Yogamani. Deep semantic segmentation for automated driving: Taxonomy, roadmap and challenges. In *2017 IEEE 20th international conference on intelligent transportation systems (ITSC)*, pages 1–8. IEEE, 2017.
- [16] C. R. Souza and P. E. Santos. Probabilistic logic reasoning about traffic scenes. In *Conference Towards Autonomous Robotic Systems*, pages 219–230. Springer, 2011.
- [17] L. von Rueden, S. Mayer, K. Beckh, B. Georgiev, S. Gieselbach, R. Heese, B. Kirsch, J. Pfommer, A. Pick, R. Ramamurthy, M. Walczak, J. Garcke, C. Bauckhage, and J. Schuecker. Informed machine learning – a taxonomy and survey of integrating knowledge into learning systems, 2020.
- [18] M. Zand, S. Doraisamy, A. A. Halin, and M. R. Mustaffa. Ontology-based semantic image segmentation using mixture models and multiple crfs. *IEEE Transactions on Image Processing*, 25(7):3233–3248, 2016.