

Zero-shot Multi-Domain Dialog State Tracking Using Prescriptive Rules

Edgar Altszyler^{*1,2}, Pablo Brusco^{*3}, Nikoletta Basiou^{†4}, John Byrnes⁵ and Dimitra Vergyri⁵

¹*Departamento de Computación, FCEyN, Universidad de Buenos Aires, Argentina*

²*Instituto de Investigación en Ciencias de la Computación, CONICET-UBA, Argentina*

³*ASAPP, USA*

⁴*Amazon, Alexa-AI*

⁵*SRI International, USA*

Abstract

In this work, we present a framework for incorporating declarative logical rules in state-of-the-art neural networks, enabling them to learn how to handle unseen labels without the introduction of any new training data. The rules are integrated into existing networks without modifying their architecture, through an additional term in the network's loss function that penalizes states of the network that do not obey the designed rules. As a case study, the framework is applied to an existing neural-based *Dialog State Tracker*. Our experiments demonstrate that the inclusion of logical rules allows the prediction of unseen labels, without deteriorating the predictive capacity of the original system.

Keywords

Zero-shot Learning, Prescriptive rules, Neural network, Natural Language Understanding

1. Introduction

When deploying machine-learning-based systems, it is common for users to detect problems related to functionalities that do not meet the expected requirements. In particular, in dialog systems, problems arise when for a certain input, the model makes a prediction that is different from the user's inferred decision. This is often due to the model structure and to the inherent characteristics of the dataset used to train the models. In the same direction, new user requirements for the existing system may require outputting **unseen labels** (not present in training data). In such situations, a typical solution consists of the collection of new annotated data aligned with the expected functionality. However, collecting new data every time such a need arises is an expensive and time-consuming effort, therefore, an alternative approach is desired.

In this work, we propose a solution to the above mentioned problems by incorporating **prescriptive logical rules** into learned neural network models. These rules are designed by

15th International Workshop on Neural-Symbolic Learning and Reasoning


**These authors contributed equally to this work. This work was performed while the authors were still at SRI.*

† Work done prior to joining Amazon.

✉ ealtszyler@dc.uba.ar (E. Altszyler^{*}); pbrusco@dc.uba.ar (P. Brusco^{*}); nbbasiou@amazon.com (N. Basiou[†]); john.byrnes@sri.com (J. Byrnes); dimitra.vergyri@sri.com (D. Vergyri)



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

domain experts and can influence the system output, enabling also the prediction of unseen labels. Similar to other works [1, 2, 3, 4, 5], we use differentiable first-order logic (FOL) that has proven useful for integrating knowledge into a neural-symbolic system.

We apply our logic rules framework to Dialog State Tracking, a challenging and complex task in the field of Dialog Systems. More specifically, we extend the Multi-Domain Neural Belief State Tracker (MDNBT), proposed in [6] and recently incorporated as one of the state of the art dialog state trackers in ConvLab, an open-source multidomain end-to-end dialog system platform released under the Dialog State Tracker Challenge (DSTC8) [7].

The main contributions of our work are the following: a) we enhance a Neural-based Dialog State Tracker with logic rules, without degrading the performance of the base system, b) we show that the addition of the logic rules allow the predictions of unseen labels which can be very useful in the case of unlabeled or partially labeled data.

2. Related Work

Due to the increasing popularity of neural network models for supervised learning, there is a growing body of material related to the inclusion of structural knowledge as a tool of biasing certain models decisions and as a way to mitigate the uninterpretability of results. One way to introduce this knowledge is to integrate logical rules through the use of FOL — a declarative language that can represent high-level knowledge [1, 2, 3, 4, 5, 8, 9, 10, 11, 12, 13, inter alia]. In the majority of prior works, different types of posterior regularization terms are implemented to affect the optimization process. For example, in a seminal work by Hu et al. [1], the authors propose a general teacher-student framework approach in which the model simultaneously learns from labeled data and logical rules through an iterative process that shifts the parameters of CNN and RNN networks in the tasks of sentiment classification and named entity recognition. Although these works generally apply rules as functions of the network inputs and outputs, there are some that can predicate over the internal values of the network’s neurons [5]. These rules can be included in the existing neural network architectures to guide the training and predictions without additional learning parameters [1, 11].

There are also works that use these techniques in semi-supervised or even zero-shot learning settings. For example, in [1, 3] it is shown that the rules can be applied over unlabeled data. Furthermore, Donadello and Serafini [14] showed that the use of rules in a semantic image interpretation task allows predicting unseen visual relationships. However, to our knowledge, this type of framework has not been applied to the fully-unsupervised problem of **unseen-labels prediction**.

3. Our approach

The proposed framework consists of the addition of a plug-in component into an existing computational graph with **no extra learning parameters**. A rules-dependent loss term is introduced into the system’s loss function as a way of integrating rules to the learning process and as a way of allowing the use of off-the-shelf optimizers. By including rules as part of the training process, we generate extra cost when the network does not satisfy a certain rule.

3.1. Neural Belief State Tracker

A *Dialog State Tracker (DST)* is a key component in task-based spoken dialog systems. It models the user’s intent at any point of an ongoing conversation [15] which is then used by the downstream dialog management component to choose the next system response. DST models estimate the *belief state*, which is the system’s internal probability distribution over possible dialog states, by taking into account the user goals at every turn as extracted by a Spoken Language Understanding (SLU) component. The dialog states are defined by a domain-specific ontology that lists the slot-value pairs that describe the constraints the users can express (e.g. price range-expensive, price range-cheap, area-west, area-east, etc) [16].

In our work, we use the Multi-domain Neural Belief State Tracker (MDNBT) which jointly identifies the domain and tracks the belief states corresponding to that domain by utilizing the semantic similarity between dialog utterances and ontology terms [6]¹. MDNBT is implemented as multi-layer networks with Bi-LSTMs to model the user and system utterance and with RNNs with a memory cell to model the flow of the conversation. We use the same set of parameters as Ramadan et al. [6].

For our experiments, we used the MultiWOZ 2.0 dataset [17], following Ramadan et al. [6]. This dataset contains 2480 single-domain dialogues and 7375 multiple-domain dialogues, in which at least two domains are involved throughout each conversation. The ontology contains a total of 663 domain-slot-values triples, distributed across 27 slots in 5 domains (restaurant, hotel, attraction, train, taxi). Here we show an example of two turns inside a conversation with their corresponding state label.

1. **utterance:** Hi, can you help me find a place to eat on the northside?
{True state: restaurant-area-north}
system: Yes, I have 15 options, do you have any preferences for the price range?
2. **utterance:** Yes, I would like an expensive restaurant
{True state: restaurant-area-north, restaurant-pricerange-expensive}
system: There are 2 expensive places, an Italian restaurant and a gastropub.

3.2. Rules definition

Rules are defined as formulas in a relaxation of FOL that represent truth values in a continuous domain (FOL Fuzzy Logic) in which the satisfaction of rules is a differentiable function that can be maximized to perform learning [2, 11]. A formula’s truthiness is represented as a real number that indicate the degree of truth or falsity of a relation defined over entities of the system².

We formulated two types of rules for this study. The type R_1 , which triggers when some specific keyword (in this case *expensive*) is explicitly uttered in a specific domain (in this case *hotel*), and the type R_2 in order to preserve belief-state predictions related to the price-range slots throughout the turns for the cases when the user’s price-range intent does not change

¹<https://github.com/osmanio2/multi-domain-belief-tracking>

²In our experiments, we are not interested in FOL functions and quantifiers, and we leave them out of the discussion.

throughout the turns. Examples of the two types of rules for the *HOTEL* domain are given below:

R_1 : **IF** the **user's utterance** contains a word like *EXPENSIVE* **AND** also contains a word like *HOTEL* **THEN** the **prediction** in the domain *HOTEL*, slot *PRICERANGE* should be *EXPENSIVE*

R_2 : **IF** the **previous prediction** for the domain *HOTEL*, slot *PRICERANGE* was *EXPENSIVE* **AND** the user did **NOT uttered** a word like *MODERATE* **AND** did **NOT uttered** a word like *CHEAP* **THEN** the next **prediction** in the domain *HOTEL*, slot *PRICERANGE* should be *EXPENSIVE*.

In these examples, underlined words represent *predicates*; bold-uppercase terms represent *logical connectives*; bold-lowercase terms represent nodes in the computational graph on the network (or concepts that are mapped into embeddings such as in the case of the user's utterance); finally, italic-uppercase words refer to the *constants* of our system.

As we will discuss in further detail in the next sections, rules are implemented in the model's computational graph through the addition of new operations that are applied over existing nodes. When the graph is evaluated for a specific instance, each rule produces a number that indicates the *truthiness* of that rule for the specific instance under evaluation. The loss function will be a function of the *truthiness*, allowing the network to learn the rules.

3.2.1. Learning mechanism

Unsatisfied rules generate a cost that the optimizer minimizes in conjunction with the misclassification cost. Our system is trained not only to learn from labels (by minimizing the cross-entropy loss function) but also to learn how to make the set of rules \mathcal{R} as true as possible (a concept called *best satisfiability* as presented in [18]). For this, we use the simplest posterior regularization approach, in which the objective function of the rule-based model is the sum of the loss function of the base MDNBT model (\mathcal{L}_{MDNBT}) and the rules' loss function (\mathcal{L}_{rules}) [11],

$$\mathcal{L}(\theta; \mathcal{D}; \mathcal{Y}; \mathcal{R}) = \mathcal{L}_{MDNBT}(\theta; \mathcal{D}; \mathcal{Y}) + w \mathcal{L}_{rules}(\theta; \mathcal{D}; \mathcal{R})$$

where θ represents the model parameters (weights and biases), \mathcal{D} , \mathcal{Y} refers to the dataset and its labels respectively, \mathcal{R} represents the set of rules of our system, w is a weighting hyperparameter that we call *rules' weight*, and \mathcal{L}_{rules} is defined as the sum of the loss of each individual rule:

$$\begin{aligned} \mathcal{L}_{rules}(\theta; \mathcal{D}; \mathcal{R}) &= \sum_{r \in \mathcal{R}} \mathcal{L}_r(\theta; \mathcal{D}) \\ &= \sum_{r \in \mathcal{R}} 1 - truthiness_r(\theta; \mathcal{D}) \end{aligned}$$

Here, each individual rule loss (\mathcal{L}_r) is defined in terms of the *truthiness_r*, i.e. the degree of truth of the rule. Finally, the optimization problem consists of finding the optimal weights and biases for the network given the composed loss function,

$$\theta^* = \arg \min_{\theta} \mathcal{L}(\theta; \mathcal{D}; \mathcal{Y}; \mathcal{R}) \quad (1)$$

The backpropagation mechanism computes gradients that update all trainable parameters in the network that are reachable from the loss function. For weights and biases to be reachable, the operations that define the loss function have to be fully differentiable. Thus the logical operations that define truthiness of the rules need to be differentiable.

3.2.2. Formulas and Predicates

We define two types of formulas: (i) atomic formulas: predicates applied to constants and nodes of the computational graph; and (ii) composed formulas, which are built up from atomic formulas using the Boolean connectives.

Predicates define relations among entities of the neural network. For example, in the case of the MDNBT model, we can refer to the embedding representation of a word in the input data; to the actual belief states (b_t) – a slot-specific distribution of probabilities estimated by the MDNBT output layer [6]; to previous belief states (b_{t-1}); etc. Also, predicates may refer to external values to the computational graph (i.e constants) such as a pre-trained word embedding for the word *HOTEL*. When the computational graph is evaluated, each predicate returns a truthiness value. That is, a value between 0 and 1, with 1 being the highest confidence in the truth of a predicate.

The implementation of a predicate, as opposed to that of a logic connective, strongly depends on the underlying architecture and on the input representations. For example, the predicate contains a word like *CHEAP* can be expressed as the cosine similarity function over nodes of the computational graph: $In_{tthresh=t}(utt, \text{CHEAP})$. This function may check if there is any word in the utterance whose cosine similarity with the word embedding for *cheap* is greater than threshold t .

3.2.3. Logical operators

Logic operators (namely $\wedge, \vee, \neg, \rightarrow$) are implemented through the following equations based on *Product* t -norm for conjunction, *s*-norm for disjunction, and the *residuum* of the t -norm for the implication (see [19] and [11] for further details).

$$\begin{aligned}\neg X &\stackrel{\text{def}}{=} 1 - X, \\ X \wedge Y &\stackrel{\text{def}}{=} X \cdot Y, \\ X \vee Y &\stackrel{\text{def}}{=} \neg(\neg X \wedge \neg Y) = X + Y - XY \\ X \rightarrow Y &\stackrel{\text{def}}{=} \neg(X \wedge \neg Y) = 1 - X(1 - Y)\end{aligned}$$

Having defined all the aforementioned components, we can represent rules in terms of logic formulas. For example, R_1 is:

$$R_1 \equiv (In_{tthresh=t}(utt, \text{EXPENSIVE}) \wedge In_{tthresh=t}(utt, \text{HOTEL})) \rightarrow \text{Assert}(b_t, \text{HOTEL-PRICERANGE-EXPENSIVE})$$

where $\text{Assert}(b_t, \text{HOTEL-PRICERANGE-EXPENSIVE})$ is a predicate that returns the model's belief state probability b_t in the index corresponding to the *HOTEL-PRICERANGE-EXPENSIVE* state.

3.2.4. Antecedent and Consequent learning

When a gradient-based method is used to solve the optimization problem of Eq. 1, the parameters θ of the network are updated in the opposite direction to the gradient of \mathcal{L} . The rule loss term produces an update associated with it. For example, in the case of a rule with the form $r = X \rightarrow Y$ in which the loss is $\mathcal{L}_r = X(1 - Y)$, the update associated to the rule looks like:

$$\Delta_r \theta = -\lambda(d_X(\mathcal{L}_r)d_\theta(X)) + d_Y(\mathcal{L}_r)d_\theta(Y))$$

where λ is the learning rate and the partial derivatives of the implication are $d_X(\mathcal{L}_r) = 1 - Y$ and $d_Y(\mathcal{L}_r) = -X$.

For example, in the case where this implication is not satisfied ($X = 1$ and $Y = 0$), the partial derivatives are $d_X(\mathcal{L}_r) = 1$ and $d_Y(\mathcal{L}_r) = -1$, and the update is,

$$\Delta_r \theta = -\lambda d_\theta(X) + \lambda d_\theta(Y)$$

The network will update θ in the direction of growth of Y and in the direction of decrease of X . That is, during the learning step the antecedent tends to decrease and the consequent tends to increase simultaneously.

Depending on the rule, one may want the antecedent or the consequent learning to be *frozen* (i.e. that the learning process occurs only through one of them). This is the case of some *if-type* rules like R_2 , in which we are not interested in learning how to make the *condition True*, but we want to make the *then branch True* (in case the condition is met). For these types of cases, one can use two different alternatives: either implementing predicates using non-derivable functions such as *argmax*; or by programmatically stopping the back-propagation in the corresponding subterms.

4. Results and Discussion

The main question we address in this section is: *can we extend a system with a new slot without any additional data and without degrading the existing system?*

To answer this question, we simulated this scenario by removing all existing annotations for the *PRICERANGE* slot in the training set of the *MULTIWOZ* dataset. Next, we built a set of twelve rules designed to learn the *PRICERANGE* slot values. Six of these rules were in the form of R_1 , and six were in the form of R_2 (as described in Section 3.2), addressing the six possible combinations of domains (hotel and restaurant) and price ranges (cheap, moderate and expensive).

We compare our models to a *base* MDNBT model that does not contain rules and in which we removed the *PRICERANGE* slot-value pairs from the ontology. In our experiments, we used 100 Bi-LSTM cells and we trained the models from scratch using the ADAM optimizer [20] with batch size 64 for 600 epochs. A dropout rate [21] of 50% was used in all the intermediate representations. Also, all the weights were initialized using normal distribution of zero mean and unit variance and biases were initialized to zero. For the rule-based MDNBT we trained models with four different values of the rules' weight, i.e., $w = 10, 30, 100, 300$.




Figure 1: Performance of rules-based MDNBT models for the *PRICERANGE* slots (left), and for the rest of the slots (right). Each star shows the F_1 for 3 different runs per weight, and the mean values are shown as circles. The mean F_1 , in the *PRICERANGE* slots, is 0.093, 0.493, 0.521 and 0.436 for the values of w of 10, 30, 100, 300 respectively. While the mean F_1 in the other slots is 0.823, 0.697, 0.832, and 0.581 for the w of 10, 30, 100, and 300 respectively. Performance of the *base MDNBT model* is also included in the right plot, which is computed as the mean among 6 runs ($\mu = 0.817 \pm 0.085$).

The F_1 performance of the rules-based MDNBT models for the *PRICERANGE* slot and for all the remaining slots are depicted in Fig. 1. F_1 is measured by considering the correct and incorrect predictions in each slot of each domain. The performance of the base MDNBT model for the remaining slots is also shown (right plot). From the left panel, it is evident that the rule-based MDNBT model shows predictive capability over the *PRICERANGE* slot values without any training data (zero-shot learning). The increase in the performance measured in the *PRICERANGE* slots for different rules' weight values affects the overall performance on the rest of the slots in the ontology depending on the rules' weight value. As can be seen from the right plot, for lower rules' weight values (e.g. $w = 10$) in the rule-based MDNBT there is not an appreciable performance drop compared to the base MDNBT (0.7% relative decrease), while for large rules' weights (e.g. $w = 300$) there is a considerable degradation of performance (i.e., 29% relative decrease). A selection of an optimal rules' weight (e.g. $w = 30$) can guarantee the optimal trade-off between the performance on unlabeled data (i.e. *PRICERANGE* slot) and the performance on labeled data (e.g. the remaining slots in the ontology).

Our results do not show a significant performance degradation in the other slots (right panel) for low rules-weight values ($w = 10, 30$), since it does not show significant difference with the base model (two-sided t -test p -val > 0.1 in both w values). However, using $w = 100$ or $w = 300$ we see a notable decrease in the general performance (p -val = 0.06 and p -val = 0.004 respectively with two-sided t -tests). That is, the system learns how to identify price ranges at the expense of producing unwanted effects in the performance of the rest of the slots.

From the experimental results, we observe that it is possible to integrate rules into an existing system to allow the prediction of unseen labels without degrading the predictive capabilities

over the rest of the labels (as it is the case with $w = 30$). However, it is necessary to pay special attention to the trade-off that is generated between learning the rules and the degradation of the system. In particular, it is important to notice that the weights depend on the number of times the rules are actually satisfied, the number of rules, and the design properties of the system.

5. Scope and limitations

In this work we show how a set of rules can be incorporated into a Neural Network to predict a new category that did not exist in the training set.

It is worth noting that this rule-based setup completely depends on the coverage of a set of rules established by the domain experts, and also, that the rules/weight/predicate and values/slot names are problem specific.

The presented approach serves specifically for the case in which there are new requirements for a system that can easily be covered – allowing the creation of an updated version of the system without the need of additional annotations.

6. Future Work

As we described in section 2, there are other methods to integrate rules in neural networks, and also, there are new state-of-the-art Dialogue State Tracking models that produce better metrics [22]. Since the objective of this work was to show that it is possible to perform Zero-Shot Learning by adding rules to an existing neural network, we have chosen a simple and transparent rule integration method, and we have not focused on the selection of the State Tracking model. In the future, we plan to compare different methods of integration of rules in state-of-the-art models.

Campagna et al. [23] have used data synthesis techniques for transferring knowledge into new domains (Zero-Shot Transfer Learning). In the future, we plan to adapt this method for our task (i.e. predict slots not seen during training) to compare with our approach. We believe that the integration of rules in neural networks has the benefit of using existing information in the internal states and to predicate, for example, about the prediction of the model during training.

7. Conclusions

This paper presented how the addition of prescriptive logical rules designed by domain experts can enable neural networks to predict unseen labels without the need for creating new labeled training data. The rules are integrated into an existing neural network without modifying the original architecture. A posterior regularization approach was used to introduce the rules into the learning process, penalizing the objective function when inputs and the internal state of the network do not obey one of the designed rules. Our rules-based framework was applied and tested to an existing neural-based *Dialog State Tracker* for Dialog Systems where rules were implemented so that the model learns to identify *PRICERANGE* labels, which were not seen during training. Our experiments showed that the inclusion of logical rules allows the prediction of new labels, without jeopardizing the predictive capacity on the rest of the data. It

is finally worth noting that our rules-based solution is independent of the neural network model and thus can be applied to any application (and neural network model) given the formulation of appropriate rules.

References

- [1] Z. Hu, X. Ma, Z. Liu, E. Hovy, E. Xing, Harnessing deep neural networks with logic rules, in: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Berlin, Germany, 2016, pp. 2410–2420. URL: <https://www.aclweb.org/anthology/P16-1228>. doi:10.18653/v1/P16-1228.
- [2] K. Sikka, A. Silberfarb, J. Byrnes, I. Sur, E. Chow, A. Divakaran, R. Rohwer, Deep adaptive semantic logic (dasl): Compiling declarative knowledge into deep neural networks, arXiv preprint arXiv:2003.07344 (2020).
- [3] J. Xu, Z. Zhang, T. Friedman, Y. Liang, G. Van Den Broeck, A semantic loss function for deep learning with symbolic knowledge, 35th International Conference on Machine Learning, ICML 2018 12 (2018) 8752–8760. arXiv:arXiv:1711.11157v2.
- [4] M. Fischer, M. Balunovic, D. Drachler-Cohen, T. Gehr, C. Zhang, M. Vechev, DL2: Training and querying neural networks with logic, in: International Conference on Machine Learning, 2019, pp. 1931–1941.
- [5] T. Li, V. Srikumar, Augmenting neural networks with first-order logic, arXiv preprint arXiv:1906.06298 (2019). doi:10.18653/v1/P19-1028.
- [6] O. Ramadan, P. Budzianowski, M. Gasic, Large-scale multi-domain belief tracking with knowledge sharing, in: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), 2018, pp. 432–437. doi:10.18653/v1/P18-2069.
- [7] S. Lee, Q. Zhu, R. Takanobu, X. Li, Y. Zhang, Z. Zhang, J. Li, B. Peng, X. Li, M. Huang, J. Gao, Convlab: Multi-domain end-to-end dialog system platform, in: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 2019. doi:10.18653/v1/P19-3011.
- [8] B. Zhang, X. Xu, X. Li, X. Chen, Y. Ye, Z. Wang, Sentiment analysis through critic learning for optimizing convolutional neural networks with rules, *Neurocomputing* 356 (2019) 21–30. doi:10.1016/j.neucom.2019.04.038.
- [9] G. Marra, F. Giannini, M. Diligenti, M. Gori, Integrating learning and reasoning with deep logic models, arXiv preprint arXiv:1901.04195 (2019). doi:10.1007/978-3-030-46147-8_31.
- [10] B. Chen, Z. Hao, X. Cai, R. Cai, W. Wen, J. Zhu, G. Xie, Embedding Logic Rules Into Recurrent Neural Networks, *IEEE Access* 7 (2019) 14938–14946. doi:10.1109/ACCESS.2019.2892140.
- [11] E. van Krieken, E. Acar, F. van Harmelen, Analyzing differentiable fuzzy logic operators, arXiv preprint arXiv:2002.06100 (2020).
- [12] G. Marra, M. Diligenti, F. Giannini, M. Gori, M. Maggini, Relational neural machines, arXiv preprint arXiv:2002.02193 (2020).

- [13] M. Diligenti, M. Gori, C. Sacca, Semantic-based regularization for learning and inference, *Artificial Intelligence* 244 (2017) 143–165.
- [14] I. Donadello, L. Serafini, Compensating supervision incompleteness with prior knowledge in semantic image interpretation, in: *2019 International Joint Conference on Neural Networks (IJCNN)*, IEEE, 2019, pp. 1–8.
- [15] S. Young, Cognitive user interfaces, *IEEE Signal Processing Magazine* (2010). doi:10.1109/MSP.2010.935874.
- [16] N. Mrkšić, I. Vulić, Fully statistical neural belief tracking, *arXiv preprint arXiv:1805.11350* (2018).
- [17] P. Budzianowski, T.-H. Wen, B.-H. Tseng, I. Casanueva, S. Ultes, O. Ramadan, M. Gašić, Multiwoz-a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling, in: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018. doi:10.18653/v1/D18-1547.
- [18] I. Donadello, L. Serafini, A. d’Avila Garcez, Logic tensor networks for semantic image interpretation, in: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17, 2017*, pp. 1596–1602. URL: <https://doi.org/10.24963/ijcai.2017/221>. doi:10.24963/ijcai.2017/221.
- [19] L. Serafini, A. D. Garcez, Logic tensor networks: Deep learning and logical reasoning from data and knowledge, *CEUR Workshop Proceedings* 1768 (2016). arXiv:arXiv:1606.04422v2.
- [20] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: *Proceedings of ICLR*, 2014.
- [21] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: a simple way to prevent neural networks from overfitting, *The journal of machine learning research* 15 (2014) 1929–1958.
- [22] H. Lee, J. Lee, T.-Y. Kim, Sumbt: Slot-utterance matching for universal and scalable belief tracking, in: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 5478–5483. doi:10.18653/v1/P19-1546.
- [23] G. Campagna, A. Foryciarz, M. Moradshahi, M. Lam, Zero-shot transfer learning with synthesized data for multi-domain dialogue state tracking, in: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 122–132. doi:10.18653/v1/2020.acl-main.12.