

---

# Message Passing Query Embedding

---

Daniel Daza<sup>1 2 3</sup> Michael Cochez<sup>1 3</sup>

## Abstract

Recent works on representation learning for Knowledge Graphs have moved beyond the problem of link prediction, to answering queries of an arbitrary structure. Existing methods are based on ad-hoc mechanisms that require training with a diverse set of query structures. We propose a more general architecture that employs a graph neural network to encode a graph representation of the query, where nodes correspond to entities and variables. The generality of our method allows it to encode a more diverse set of query types in comparison to previous work. Our method shows competitive performance against previous models for complex queries, and in contrast with these models, it can answer complex queries when trained for link prediction only. We show that the model learns entity embeddings that capture the notion of entity type without explicit supervision.

## 1. Introduction

Knowledge Graphs (KG) are useful data structures for encoding information from different domains, by representing entities and relations of different types between them. Tasks of interest that can be addressed with knowledge graphs include information retrieval (Dalton et al., 2014), question answering (Vakulenko et al., 2019; Huang et al., 2019), and natural language processing (Logan et al., 2019). A common way to answer a question using a KG is to pose it as a structured query (for example, using the SPARQL query language (Harris et al., 2013)). The query is then answered via logical inference, using the information present in the graph. However, knowledge graphs are usually incomplete, either due to the construction process, or their dynamic nature. This means that there will be cases where these systems return *no answer* for a query.

To address this problem, we follow recent works that pro-

<sup>1</sup>Vrije Universiteit Amsterdam <sup>2</sup>University of Amsterdam  
<sup>3</sup>Discovery Lab, Elsevier, The Netherlands. Correspondence to: Daniel Daza <d.dazacruz@vu.nl>.




Figure 1. Message Passing Query Embedding takes as input a query graph and outputs a query embedding. Features for each node in the query graph are embeddings of entities in the KG, or type embeddings. A GNN propagates information across the graph, and an aggregation function yields the query embedding.

pose to map the query and all entities in the KG to an embedding space (Hamilton et al., 2018; Wang et al., 2018; Mai et al., 2019), where we can compute similarity scores to produce a ranked list of answers. We propose Message Passing Query Embedding (MPQE), motivated by the observation that queries over a KG can be represented by small graphs, where nodes correspond to constant entities and variables. We employ a Graph Neural Network (GNN) to perform message passing on the query graph, and an aggregation function to combine all the messages in a single vector, which acts as a representation of the query in the embedding space. By training on the task of query answering, our method learns jointly an embedding for each entity in the KG, and type embeddings for variables.

Our contributions can be summarized as follows: 1) We propose a novel method to embed queries over knowledge graphs, that addresses limitations of previous works in terms of computational complexity, and the diversity of query structures that it admits; 2) we show experimentally that MPQE is competitive with previous work on complex query embedding across multiple query structures. Furthermore, we provide evidence of the superior generalization of MPQE by training for link prediction only, and testing on complex queries. While previous models fail in this setting, we show that MPQE retains predictive performance on complex queries; 3) we conduct a qualitative analysis of the entity embeddings produced by query embedding methods, and show that MPQE learns in an unsupervised way a structured embedding space where entities cluster according to type.

## 2. Problem Definition

We define a Knowledge Graph (KG) as a tuple  $(\mathcal{V}, \mathcal{E}, \mathcal{R}, \mathcal{T})$ , where  $\mathcal{V}$  is a set of nodes representing entities, and  $\mathcal{E}$  a set of typed edges between the nodes. A function  $\tau : \mathcal{V} \rightarrow \mathcal{T}$  assigns a type to every node, where  $\mathcal{T}$  is a set of entity types. Each edge in  $\mathcal{E}$  corresponds to a relation between two nodes  $v_i$  and  $v_j \in \mathcal{V}$ , that we denote by  $r(v_i, v_j)$ , where  $r \in \mathcal{R}$  is a relation type.

Given a KG, we can pose queries that seek for an entity satisfying certain conditions. We consider queries in *conjunctive form*, formed by a conjunction of binary predicates where the arguments are entities or variables. To illustrate this, consider a KG of an academic institution where researchers work on topics, and topics are related to projects, and the following query: “select all projects  $P$ , such that topic  $T$  is related to  $P$ , and both *alice* and *bob* work on  $T$ .” This query asks for entities  $P$  that satisfy the following condition:

$$P.\exists T, P : \text{related}(P, T) \wedge \text{works\_on}(\text{alice}, T) \wedge \text{works\_on}(\text{bob}, T).$$

In general, a query  $q$  is defined by a condition on a *target variable*  $V_t$  as follows:

$$q = V_t.\exists V_1, \dots, V_m : r_1(a_1, b_1) \wedge \dots \wedge r_m(a_m, b_m), \quad (1)$$

where  $r_i \in \mathcal{R}$ , and  $a_i$  and  $b_i$  are either entities in the KG, or query variables in  $\{V_t, V_1, \dots, V_m\}$ . An entity is a correct answer if it satisfies the condition defined by the query. We can address the problem of returning a list of entities that satisfy the query, by mapping the query to an embedding and computing its similarity with the embeddings of entities in the KG (Hamilton et al., 2018). Formally, we optimize an embedding  $\mathbf{e}_v \in \mathbb{R}^d$  for every entity  $v \in \mathcal{V}$ , and we define an embedding method for the query that maps it to a vector  $\mathbf{q} \in \mathbb{R}^d$ . We then compute the cosine similarity between  $\mathbf{q}$  and an entity embedding  $\mathbf{e}_v$  as follows:

$$\text{score}(\mathbf{q}, \mathbf{e}_v) = \frac{\mathbf{q}^\top \mathbf{e}_v}{\|\mathbf{q}\| \|\mathbf{e}_v\|}. \quad (2)$$

This score determines the rank of an entity as a possible answer to the query.

## 3. Related Work

Multiple approaches for machine learning on graphs consider embedding the graph into a vector space via link prediction (Bordes et al., 2013; Wang et al., 2014; Yang et al., 2015). The applicability of these methods for answering complex queries is limited. For each link that needs to be predicted, these methods must consider all possible entities, which is exponential in the size of the query. Our method is based on an architecture that directly encodes the query

into an embedding, which provides our method with a linear complexity in the size of the graph. Recent works have also addressed the problem of embedding a query to retrieve approximate answers, by partitioning the query graph in different subgraphs, so that candidate answers can be provided for each of them (Zhang et al., 2018). In (Wang et al., 2018), the authors pre-train embeddings using an algorithm inspired by TransE. Instead of relying on a separate pre-training step, we learn with an objective that optimizes entity embeddings for the task of query answering directly. The most related approaches to our work consider embedding queries directly in the embedding space (Hamilton et al., 2018; Mai et al., 2019; Ren et al., 2020), by applying a sequence of projection and intersection operators that follow the structure of the query graph. These methods are constrained to having entities only at the leaves of the Directed Acyclic Graph (DAG) defined by the query. Furthermore, the use of projection and intersection mechanisms requires training with multiple query structures that comprise both chains and intersections. Our method has a more general formulation that enables it to i) encode a general set of query graphs, without constraints on the location of entities in the query, and ii) learn from single-hop link prediction training alone, and still generalize to larger queries.

## 4. Message Passing Query Embedding

As noted in previous work (Hamilton et al., 2018; Wang et al., 2018), queries in conjunctive form can be represented as a DAG. In this graph, the leaf nodes correspond to constant entities, the root to the variable to be retrieved, and any intermediate nodes to other variables in the query. Given a query of the form given in eq. 1, we define the *query graph* as the tuple  $(\mathcal{V}_q, \mathcal{E}_q, \mathcal{R}_q)$ . Here,  $\mathcal{V}_q$  is the union of nodes for constant entities  $\mathcal{V}_{qe}$  and variables  $\mathcal{V}_{qv}$  in the query. To construct the set of edges  $\mathcal{E}_q$ , we add one edge  $r_i(a, b)$  for each predicate in the query.

### 4.1. Model Definition

Message Passing Query Embedding (MPQE) learns a matrix of *entity embeddings*  $\mathbf{M}_e \in \mathbb{R}^{|\mathcal{V}| \times d}$ , where each row contains an embedding for each entity in the KG, and  $d$  is the dimension of the embedding space. Additionally, it also learns a matrix of *type embeddings*  $\mathbf{M}_t \in \mathbb{R}^{|\mathcal{T}| \times d}$ , with one embedding for each type of entity in the KG. Assume we define an ordering on the sets  $\mathcal{E}$  and  $\mathcal{T}$ . The function  $\phi_e(v)$  returns the row of  $\mathbf{M}_e$  corresponding to  $v$ , for all entities  $v \in \mathcal{E}$ . Similarly, for all types  $v \in \mathcal{T}$ ,  $\phi_t(v)$  returns its corresponding row of  $\mathbf{M}_t$ .

Given a query graph, we start by initializing each of the nodes with a vector representation  $\mathbf{h}_v^{(0)}$ , from  $\mathbf{M}_e$  if it corresponds to a constant entity, or from  $\mathbf{M}_t$  for a variable.

Formally,  $\mathbf{h}_v^{(0)} = \phi_e(v)$  if  $v \in \mathcal{V}_{qe}$ , and  $\mathbf{h}_v^{(0)} = \phi_t(v)$  if  $v \in \mathcal{V}_{qv}$ . We proceed by applying  $L$  steps of *message passing* with a Relational Graph Convolutional Network (R-GCN) (Schlichtkrull et al., 2018), which updates the representation of a node taking into account its neighbors and the type of the relations involved. The representation for node  $v$  at step  $l + 1$  in the R-GCN is defined as follows:

$$\mathbf{h}_v^{(l+1)} = f \left( \mathbf{W}_0^{(l)} \mathbf{h}_v^{(l)} + \sum_{r \in \mathcal{R}} \sum_{j \in \mathcal{N}_v^r} \frac{1}{|\mathcal{N}_v^r|} \mathbf{W}_r^{(l)} \mathbf{h}_j^{(l)} \right), \quad (3)$$

where  $f$  is a non-linearity,  $\mathcal{N}_v^r$  is the set of neighbors of node  $v$  through relation type  $r$ , and  $\mathbf{W}_0^{(l)}$  and  $\mathbf{W}_r^{(l)}$  are parameters of the model. After  $L$  applications of an R-GCN layer, the node representations are combined into a single vector that acts as the embedding of the query, by means of an *aggregation* function  $\phi$ :

$$\mathbf{q}_\phi = \phi \left( \{\mathbf{h}_v^{(L)} | v \in \mathcal{V}_{qe} \cup \mathcal{V}_{qv}\} \right). \quad (4)$$

**Fixed message passing** We consider as aggregation functions sum and max pooling over all  $\mathbf{h}_v^{(L)}$ . We also experiment with the concatenation of representations from hidden layers of the R-GCN, followed by an MLP and a sum over nodes in the query graph. We denote this function as CMLP:

$$\mathbf{q}_{\text{CMLP}} = \sum_{v \in \mathcal{V}_q} \text{MLP} \left( [\mathbf{h}_v^{(1)}, \dots, \mathbf{h}_v^{(L)}] \right). \quad (5)$$

**Dynamic embedding** Let  $D$  denote the *diameter* of the query graph (the longest shortest path between two nodes in the graph). We propose a dynamic query embedding method, by noting that at most  $D$  message passing steps are required to propagate messages from all nodes, to the target node. The method performs  $D$  steps of message passing, and it then selects the representation  $\mathbf{h}_{v_t}^{(D)}$  of the target node  $v_t$  as the embedding of the query. We denote this as the Target Message (TM) function.

**Training** Following previous work on query embedding (Hamilton et al., 2018), we optimize all parameters using gradient descent on a contrastive loss function, where given a query  $q$  and its embedding  $\mathbf{q}$ , a positive sample  $v^+$  corresponds to an entity in the knowledge graph that answers the query, and a negative sample  $v^-$  is an entity sampled at random, that is not an answer to the query but has the correct type. We minimize a margin loss function:

$$\mathcal{L}(q) = \max(0, 1 - \text{score}(\mathbf{q}, \mathbf{h}_{v^+}^{(0)}) + \text{score}(\mathbf{q}, \mathbf{h}_{v^-}^{(0)})). \quad (6)$$

## 5. Experiments

We evaluate the performance of MPQE in query answering over knowledge graphs, by considering 7 different query

structures (detailed in Appendix A). All the code to reproduce our experiments is available online <sup>1</sup>.

**Datasets** We use publicly available knowledge graphs that have been used in the literature of graph representation learning and query answering (Ristoski et al., 2016; Hamilton et al., 2018) containing from thousands to millions of entities – **AIFB**: a KG of an academic institution, where entities are persons, organizations, projects, publications, and topics; **MUTAG**: a KG of carcinogenic molecules, where entities are atoms, bonds, compounds, and structures; **AM**: contains relations between different artifacts in the Amsterdam Museum, including locations, makers, and others; **Bio**: a dataset of a biological interaction network containing entities of type drug, disease, protein, side effect, and biological processes. Their statistics can be found in Appendix B.

**Query generation** We follow the evaluation procedure of Hamilton et al. (2018): to obtain query graphs, we sample subgraphs from the KG. Each subgraph specifies the entities and the types of variables in the query, and the correct answer to the query. For each query we also obtain a negative sample, and in the case of query graphs with intersections, a *hard* negative sample. This is an entity that would be a correct answer, if the conjunction represented by the intersection is relaxed to a disjunction. Before sampling subgraphs for training, we remove edges from the KG. We then guarantee that subgraphs sampled to generate query graphs for testing rely on at least one of these remove edges.

**Experimental setup** With the exception of the TM aggregation function (where the number of message passing steps is given by the query diameter), we use 2 R-GCN layers. For aggregation functions with MLPs, we use two fully-connected layers, and in all cases we use ReLU for the nonlinearities. We minimize eq. 6 using the Adam optimizer with a learning rate of 0.01, and use an embedding dimension of 128. We train the models for 1-hop link prediction until convergence, and then on the full set of query structures. As a baseline we evaluate the Graph Query Embedding (GQE) method by Hamilton et al. (2018) with their TransE, DistMult, and Bilinear variants.

### 5.1. Results

**Query answering** We report the area under the ROC curve (AUC) and the Average Percentile Rank (APR) on the test set. The results for the query answering task are shown in Table 1. We observe that MPQE obtains competitive performance in comparison with GQE across different datasets. MPQE underperforms in the MUTAG dataset, which we identified as the dataset with the least diverse set of relations. We noticed that while GQE-DistMult handles hard

<sup>1</sup><https://github.com/dfdazac/mpqe>

Table 1. Results on query answering averaged across different query structures. Highlighted values denote the best variant within each group of GQE and MPQE models.

Method	AIFB		MUTAG		AM		Bio	
	AUC	APR	AUC	APR	AUC	APR	AUC	APR
GQE-TransE	83.1	<b>86.7</b>	78.8	81.0	80.9	82.3	87.4	88.9
GQE-DistMult	<b>83.8</b>	86.0	<b>80.6</b>	<b>81.1</b>	<b>82.9</b>	83.2	90.0	90.3
GQE-Bilinear	83.4	83.3	78.5	79.7	80.7	<b>84.4</b>	<b>90.5</b>	<b>90.8</b>
MPQE-TM	84.9	87.4	<b>76.7</b>	<b>77.6</b>	<b>84.2</b>	<b>86.3</b>	88.8	89.8
MPQE-sum	84.7	86.8	74.6	73.1	80.9	83.6	90.0	<b>90.5</b>
MPQE-max	83.4	85.9	74.9	72.7	80.9	82.5	88.3	88.7
MPQE-CMLP	<b>86.3</b>	<b>89.1</b>	74.3	72.5	82.5	85.5	<b>90.1</b>	90.2

Table 2. Generalization results on complex query answering (AUC) when training for simple link prediction. The results show the performance on a test set of queries with chains only (ch), and the complete set of queries with chains and intersections (all). Dashes indicate not better than random results.

Method	AIFB		MUTAG		AM		Bio	
	ch	all	ch	all	ch	all	ch	all
GQE-TransE	74.0	—	<b>89.4</b>	—	85.8	—	85.5	—
GQE-DistMult	72.8	—	85.4	—	82.4	—	95.9	—
GQE-Bilinear	72.7	—	89.1	—	<b>85.9</b>	—	85.8	—
MPQE-TM	<b>77.0</b>	<b>75.5</b>	86.8	<b>77.2</b>	<b>85.0</b>	<b>81.6</b>	<b>96.4</b>	<b>83.9</b>
MPQE-sum	69.8	69.6	82.8	74.0	52.5	53.9	92.4	80.0
MPQE-max	74.1	71.9	77.1	71.6	51.2	53.0	92.0	79.9
MPQE-CMLP	69.7	69.1	84.6	74.2	51.5	53.8	89.8	78.3

negative samples well (which occur only on queries with intersections), MPQE-TM has better performance on regular samples, across all query structures. We discuss further interesting properties of MPQE in Appendix C.

**Generalization** To examine the generalization properties of MPQE, we train the models on simple queries that require 1-hop link prediction, but we carry out the evaluation using the complete set of complex query structures. Unlike MPQE, in this case GQE cannot provide an answer better than random for queries with intersections, because the intersection operator is not optimized. We thus consider two evaluations modes: evaluating on queries with chain structures only, and evaluating on the complete set of query structures (where GQE is not applicable). These modes are denoted as “ch” and “all”, respectively, in Table 2. The results of MPQE are competitive when evaluating on queries with chains only, and crucially, it also generalizes well to query structures not seen during training. This encouraging results shows that MPQE implements a mechanism that does not necessarily require training on many diverse query structures to generalize well, unlike GQE.

**Clustering** We sample 200 entity embeddings for each type in the AM dataset, and visualize them using T-




Figure 2. Visualization of the entity embeddings learned by GQE (left) and MPQE (right). Each color represents an entity type.

SNE (Maaten & Hinton, 2008) in Fig. 2 for GQE-Bilinear and MPQE-TM. We observe that MPQE has learned a structured space where entities cluster according to their type, without explicit supervision. This is in stark contrast with the embeddings of GQE, where we do not observe such a clear structure. We hypothesize that the structured embedding space of MPQE contributes to the generalization from simple link prediction to complex queries.

## 6. Conclusion

We have presented MPQE, a neural architecture to encode complex queries on knowledge graphs, that jointly learns entity and type embeddings and a straightforward message passing architecture to obtain a query embedding. Our experiments show that message passing across the query graph is a powerful mechanism for query answering, that generalizes to multiple query structures even when only trained for single hop link prediction. Qualitative results show that MPQE learns a well-structured embedding space. This result motivates future research on the application of the learned embeddings to other tasks related to KGs, such as clustering, and node and graph classification. Under this new light, MPQE can be seen as a new method for unsupervised representation learning in KGs.

While our generalization experiments highlight the general formulation of our method, we further plan to evaluate its performance on queries not restricted to constants at the leaves of the query graph in future work. By being able to encode queries independent of the position of entities and variables, we could encode queries with additional information, that could be used to condition the answers on a given context. Such an application would be useful in information retrieval and recommender systems.

Our method presents limitations when evaluating on hard negative samples. Further improvements could include improving the message passing procedure via adding attention or gating mechanisms, and extensions to more expressive query embedding representations, such as boxes (Ren et al., 2020).



## Acknowledgments

This project was funded by Elsevier’s Discovery Lab. We thank Frank van Harmelen and Erman Acar for helpful discussions and comments.

## References

- Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., and Yakhnenko, O. Translating embeddings for modeling multi-relational data. In *Advances in neural information processing systems*, pp. 2787–2795, 2013.
- Dalton, J., Dietz, L., and Allan, J. Entity query feature expansion using knowledge base links. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*, pp. 365–374. ACM, 2014.
- Hamilton, W. L., Bajaj, P., Zitnik, M., Jurafsky, D., and Leskovec, J. Embedding logical queries on knowledge graphs. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada.*, pp. 2030–2041, 2018.
- Harris, S., Seaborne, A., and Prud’hommeaux, E. SPARQL 1.1 query language. *W3C recommendation*, 21(10):778, 2013.
- Huang, X., Zhang, J., Li, D., and Li, P. Knowledge graph embedding based question answering. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pp. 105–113. ACM, 2019.
- Logan, R., Liu, N. F., Peters, M. E., Gardner, M., and Singh, S. Barack’s wife hillary: Using knowledge graphs for fact-aware language modeling. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 5962–5971, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1598. URL <https://www.aclweb.org/anthology/P19-1598>.
- Maaten, L. v. d. and Hinton, G. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov): 2579–2605, 2008.
- Mai, G., Janowicz, K., Yan, B., Zhu, R., Cai, L., and Lao, N. Contextual graph attention for answering logical queries over incomplete knowledge graphs. In *Proceedings of K-CAP 2019, Nov. 19 - 21, 2019, Marina del Rey, CA, USA.*, 2019.
- Ren, H., Hu, W., and Leskovec, J. Query2box: Reasoning over knowledge graphs in vector space using box embeddings. In *International Conference on Learning Representations*, 2020.
- Ristoski, P., De Vries, G. K. D., and Paulheim, H. A collection of benchmark datasets for systematic evaluations of machine learning on the semantic web. In *International Semantic Web Conference*, pp. 186–194. Springer, 2016.
- Schlichtkrull, M. S., Kipf, T. N., Bloem, P., van den Berg, R., Titov, I., and Welling, M. Modeling relational data with graph convolutional networks. In *The Semantic Web - 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, June 3-7, 2018, Proceedings*, pp. 593–607, 2018.
- Vakulenko, S., Fernandez Garcia, J. D., Polleres, A., de Rijke, M., and Cochez, M. Message passing for complex question answering over knowledge graphs. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pp. 1431–1440. ACM, 2019.
- Wang, M., Wang, R., Liu, J., Chen, Y., Zhang, L., and Qi, G. Towards empty answers in sparql: Approximating querying with rdf embedding. In *International Semantic Web Conference*, pp. 513–529. Springer, 2018.
- Wang, Z., Zhang, J., Feng, J., and Chen, Z. Knowledge graph embedding by translating on hyperplanes. In *Twenty-Eighth AAAI conference on artificial intelligence*, 2014.
- Yang, B., Yih, W., He, X., Gao, J., and Deng, L. Embedding entities and relations for learning and inference in knowledge bases. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6575>.
- Zhang, L., Zhang, X., and Feng, Z. TrQuery: An embedding-based framework for recommending sparql queries. In *2018 IEEE 30th International Conference on Tools with Artificial Intelligence (ICTAI)*. IEEE, 2018.




Figure 3. Query structures that we consider for the evaluation of methods on query answering. Green nodes correspond to entities, and the rest are variables in the query, with blue nodes representing the target of the query.

Table 3. Statistics of the knowledge graphs that we use for training and evaluation.

	AIFB	MUTAG	AM	Bio
Entities	2,601	22,372	372,584	162,622
Entity types	6	4	5	5
Relations	39,436	81,332	1,193,402	8,045,726
Relation types	49	8	19	56

## A. Query structures

We run experiments over 7 different query structures that combine chains and intersections, as detailed in Fig. 3.

## B. Dataset statistics

We present statistics for the datasets used in our work in Table 3. To generate splits for training, test, and validation, we follow the procedure of Hamilton et al. (2018). Given a KG, we start by removing 10% of its edges. Using this incomplete graph, we extract 1 million subgraphs, containing all the query structures outlined previously. We then restore the removed edges, and extract 11,000 additional subgraphs of all structures, ensuring that they all rely in at least one of the edges that was removed to create the training set. We split this set of query graphs into two disjoint sets, containing 1,000 queries for validation, and 10,000 for testing. We use the validation set to perform early stopping during training.

## C. MPQE properties

An interesting observation from our experiments is that the message passing mechanism alone is sufficient to provide good performance for query answering, as we can see from the results for the MPQE-TM architecture. In this model, we perform a number of steps of message passing equal to the diameter of the query, and take as query embedding the

resulting feature vector at the target node. Intuitively, this allows MPQE-TM to adapt to the structure of a query so that after message passing, all information from the entity and variable nodes has reached the target node. To confirm this intuition, we evaluate the performance of MPQE as a function of the number of message passing steps, ranging from 1 to 4. The results are shown in Figure 4, for all the query structures that we have considered. We highlight the points that correspond to the diameter of the query, and we note that the results align with our intuition about the message passing mechanism. When the number of steps matches the diameter, there is a significant increase in performance, and further steps have little effect. This supports the superior generalization observed in our experiments, in comparison with GQE, and other MPQE architectures where the number of R-GCN layers was fixed.




Figure 4. Query answering performance (AUC) as a function of the number of message passing steps (implemented by layers of an R-GCN), evaluated across different query types. Dark circles corresponds to the diameter of the corresponding query. When the number of steps matches the diameter, there is a significant increase in performance, and further steps have little effect.