Probabilistic Entity Representation Model for Chain Reasoning over Knowledge Graphs

Anonymous Author(s)
Affiliation
Address
email

Abstract

Logical reasoning over knowledge graphs is a fundamental challenge that limits its application on large and incomplete databases. Current approaches employ spatial geometries such as boxes to learn query representations that encompass the answer entities and model the logical operations of projection and intersection. However, their geometry is restrictive and leads to multiple issues. Furthermore, previous works propose transformation tricks to handle unions which results in non-closure and, thus, cannot be chained in a stream. In this paper, we propose Probabilistic Entity Representation Model (PERM) to encode entities as a Multivariate Gaussian density with mean and covariance parameters to capture its semantic position and continuous decision boundary, respectively. Additionally, we also define the closed logical operations of projection, intersection ($\cap$), and union ($\cup$) that can be chained in an end-to-end objective function. On the logical query reasoning problem, we demonstrate that PERM outperforms the state-of-the-art methods on various standard KG datasets by 6.2% in HITS@3 and 12.6% in MRR. Furthermore, we evaluate PERM's competence in the COVID-19 drug-repurposing problem and show that we are able to recommend drugs with 8.2% better F1 than current methods. Finally, we demonstrate the comprehension of PERM's query answering process by observing a low-dimensional visualization of the Gaussian embeddings.

1 Introduction

Knowledge Graphs (KGs) are structured heterogeneous graphs where information is organized as triplets of entity pair and the relation between them. This organization provides a fluid schema with applications in several domains including medical research [1, 2], e-commerce [3], and web ontologies [4, 5]. Chain reasoning is a fundamental problem in KGs, which involves answering a chain of first-order existential (FOE) queries (translation, intersection, and union) using the KGs' relation paths. A myriad of queries can be answered using such logical formulation (some examples are given in Figure 1). Current approaches [6, 7, 8] in the field rely on mapping the entities and relations into a representational latent space such that the FOE queries can be reduced to mathematical operations to further retrieve relevant answer entities.

Euclidean vectors [6, 9] provide a nice mechanism to encode the semantic position of the entities by leveraging their neighborhood relations. They utilize a fixed threshold over the vector to query for answer entities (such as a k-nearest neighbor search). However, queries differ in their breadth. Certain queries would lead to a greater set of answers than others, e.g., query `Canadians` will result in a higher number of answers than query `Canadian Turing Award winners`. To capture this query behavior, spatial embeddings [7, 8, 10, 11] learn a border parameter that accounts for broadness of queries by controlling the volume of space enclosed by the query representations. However, these spatial embeddings rely on more complex geometries such as boxes [7] which do not have a closed

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

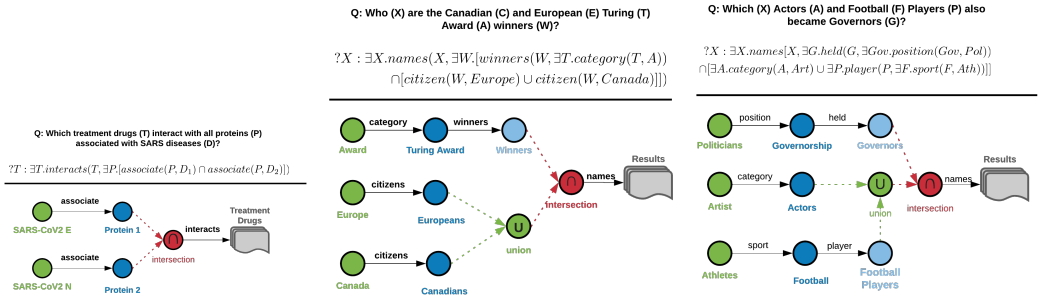

(a) Drug Repurposing (DRKG). (b) Open-domain (FB15K). (c) Organized information (DBPedia).

Figure 1: Sample FOE queries from different datasets that utilize existential quantification ($\exists$), intersection ($\cap$), and union ($\cup$) operations. The simple operations need to be chained together in an end-to-end objective to retrieve relevant results for complex queries.

solution to the union operation, e.g., the union of two boxes is not a box and thus, further FOE operations cannot be applied to a union of boxes. Additionally, their strict borders lead to ambiguity in the border case scenarios and a discontinuous distance function, e.g., a point on the border will have a much smaller distance if it is considered to be inside the box than if it is considered to be outside. This challenge also applies to other geometric enclosures such as hyperboloids [8].

Another line of work includes the use of structured geometric regions [12, 7] or densities [13, 14, 11, 15], instead of vector points, for representation learning. While these approaches utilize the representations for modeling individual entities and relations between them, we aim to provide a closed form solution to logical queries over KGs using the Gaussian density function which enables chaining the queries together. Another crucial difference in our work is handling a stream of queries. Previous approaches rely on DNF transformation which requires the entire query input. PERM does not need to rely on any transformation, since all the outputs are closed under the Gaussian space and complex queries can be consolidated in an end-to-end objective function, e.g., in Figure 2b, `Europeans` $\cup$ `Canadians` is a Gaussian mixture and the singular objective is to minimize the distance between the mixture and entity `Hinton`, whereas in the case of boxes (shown in Figure 2a), we have two independent objectives to minimize the distance from each box in the union query.

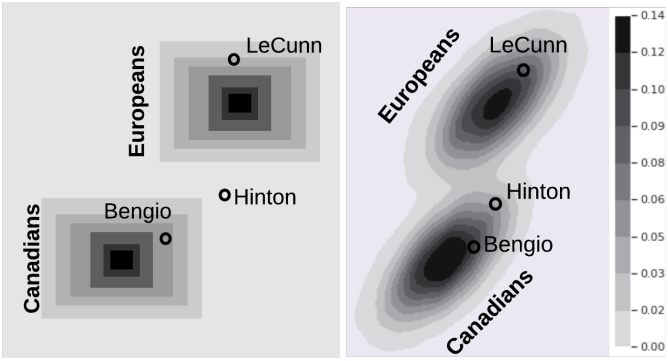

(a) Union of box queries. (b) Gaussian mixture queries.

	Entity	Bengio	LeCunn	Hinton
	Citizen	Canadian	French	British-Canadian
Union of Boxes	European	**1**	**0.211**	**1.343**
	Canadian	**0.146**	**1.569**	**1.114**
	Union	**0.146**	**0.211**	**1.114**
Gaussian Mixture	European	**1**	**0.215**	**1.347**
	Canadian	**0.146**	**1.57**	**1.118**
	Union	**0.255**	**0.628**	**0.766**

(c) Distance of entities from query space. For comparability, distances are given relative to entity `Bengio`'s distance to `European` query.

Figure 2: The figure provides the result to query `European`$\cup$`Canadian`. Entities in the darker areas have higher probability of being the answers than lighter areas. We can observe from (c) that the discontinuous borders of box geometry do not encompass the answer `Hinton`.

To alleviate the drawbacks of operations not being closed under unions and border ambiguities, we propose Probabilistic Entity Representation Model (PERM). PERM models entities as a mixture of Gaussian densities. Gaussian densities have been previously used in natural language processing [14] and graphs [15] to enable more expressive parameterization of decision boundaries. In our case, we utilize a mixture of multivariate Gaussian densities due to their intuitive closed form solution for translation, intersection, and union operations. In addition, they can also enable the use of a continuous distance function; Mahalanobis distance [16]. We utilize the mean (μ) and co-variance (Σ) parameters of multivariate Gaussian densities to encode the semantic position and spatial query area of an entity, respectively. The closed form solution for the operations allows us to solve complex queries

by chaining them in a pipeline. Figure 2 provides an example of such a case where the discontinuous boundaries of box query embeddings are not able to capture certain answers. Summarizing, our work advocates for Gaussian densities as a means for modeling KGs due to their following advantages:

1. Gaussian densities are able to provide a closed form solution to intersection and union, and also a continuous distance function. This enables us to process a chain of complex logical queries in an end-to-end objective function.
2. PERM is able to outperform the current state-of-the-art baselines on logical query reasoning over standard benchmark datasets. Additionally, it is also able to provide better drug recommendations for COVID-19 from the DRKG dataset.
3. PERM is also interpretable as Gaussian embeddings can be visualized after each query process to understand the complete query representation.

The rest of the paper is organized as follows: Section 2 presents the current work in the field. In section 3, we present PERM and define its various operations. Section 4 provides the formulation for building the reasoning chains for complex queries. We provide the experimental setup and results in section 5. We conclude our paper in section 6 and present its broader impact in section 7.

2 Related Work

A popular field related to our scope of work is multi-hop chain reasoning over KGs [17, 18, 19, 6]. These approaches utilize vector spaces to model query representation and retrieve results using a fixed threshold. While such representations are efficient at encoding semantic information, the fixed thresholds do not allow for an expressive boundary and, thus, are not best suitable for representing queries. Spatial embeddings [7, 8, 20] enhance the simple vector representations by adding a learnable border parameter that controls the spatial area around a query representation. These methods have strict borders that rely on discontinuous distance function that create ambiguity between border cases. Our model utilizes Gaussian densities and creates soft borders in terms of the variance parameter using the Mahalanobis distance. Additionally, the previous methods do not provide a closed form solution for unions which we solve using Gaussian mixture models.

Density-based embeddings have seen a recent surge of interest in various domains. Word2Gauss [14] provides a method of learning Gaussian densities for words from their distributional semantic information. In addition, the authors further apply this work to knowledge graphs [13]. Another approach [15] aims to learn Gaussian graph representations from their network connections. These methods are, however, focused on learning semantic information and do not easily extend to logical queries over knowledge graphs. PERM primarily focuses on learning spatial Gaussian densities for queries, while also capturing the semantic information. For this, we derive closed form solutions to FOE queries. In that sense, our work is also related to classic multivariate Gaussian distributions.

3 Probabilistic Entity Representation Model for Logical Operators

Knowledge Graphs (KG) $\mathcal{G} : E \times R$ are heterogeneous graphs that store entities (E) and relations (R). Each relation $r \in R$ is a Boolean function $r : E \times E \rightarrow \{True, False\}$ that indicates if the relation r exists between a pair of entities. Without loss of generality, KGs can also be organized as a set of triples $\langle e_1, r, e_2 \rangle \subseteq \mathcal{G}$, defined by the Boolean relation function $r(e_1, e_2)$. We focus on the following three FOE operations in our work; translation (t), intersection ($\cap$), and union ($\cup$). The operations are defined as below:

$$q_t[Q_t] = \{v_1, v_2, ..., v_k\} = V_t \subseteq E \; \exists \; a_1 \tag{1}$$

$$q_\cap[Q_\cap] = \{v_1, v_2, ..., v_k\} = V_\cap \subseteq E \; \exists \; a_1 \cap a_2 \cap ... \cap a_i \tag{2}$$

$$q_\cup[Q_\cup] = \{v_1, v_2, ..., v_k\} = V_\cup \subseteq E \; \exists \; a_1 \cup a_2 \cup ... \cup a_i \tag{3}$$

$$\text{where } Q_t = (e_1, r_1); \; Q_\cap, Q_\cup = \{(e_1, r_1), (e_2, r_2), ..(e_i, r_i)\} \text{ and } a_i = r_i(e_i, v_a)$$

where q_t, $q_\cap$, and $q_\cup$ are the translation, intersection, and union queries, respectively; and V_t, $V_\cap$, and $V_\cup$ are the corresponding results [10]. As we notice above, each entity has a dual nature; one as being part of a query and another as a candidate answer to a query. In PERM, we model the query nature of an entity $e_i \in E$ as a multivariate Gaussian density function; $e_i = \mathcal{N}(\mu_i, \Sigma_i)$, where the learnable

parameters μ_i (mean) and Σ_i (covariance) indicate the semantic position and the surrounding query density of the entity, respectively. As a candidate, we only consider the μ_i and ignore the Σ_i of the entity. Thus, we define the distance of a candidate entity $v_i = \mathcal{N}(\mu_i, \Sigma_i)$ from a query Gaussian $e_j = \mathcal{N}(\mu_j, \Sigma_j)$ using the Mahalanobis distance [16] given by:

$$d_{\mathcal{N}}(v_i, e_j) = (\mu_j - \mu_i)^T \Sigma_j^{-1} (\mu_j - \mu_i) \tag{4}$$

Additionally, we need to define the FOE operations for the Probabilistic Entity Representation Model. A visual interpretation of the operations; translation, intersection, and union is shown in Figure 3. The operations are defined as follows:

Translation (t). Each entity $e \in E$ and $r \in R$ are encoded as $\mathcal{N}(\mu_e, \Sigma_e)$ and $\mathcal{N}(\mu_r, \Sigma_r)$, respectively. We define the translation query representation of an entity e with relation r as q_t and the distance of resultant entity $v_t \in V_t$ from the query as d_t^q given by:

$$q_t = \mathcal{N}(\mu_e + \mu_r, (\Sigma_e^{-1} + \Sigma_r^{-1})^{-1}); \quad d_t^q = d_{\mathcal{N}}(v_t, q_t) \tag{5}$$

Intersection ($\cap$). Intuitively, the intersection of two Gaussian densities implies a random variable that belongs to both the densities. Given that the entity densities are independent of each other, we define the intersection of two entity density functions e_1, e_2 as $q_\cap$ and distance of resultant entity $v_\cap \in V_\cap$ from the query as $d_\cap^q$ given by:

$$\begin{aligned} q_\cap = \mathcal{N}(\mu_{e_1}, \Sigma_{e_1})\mathcal{N}(\mu_{e_2}, \Sigma_{e_2}) = \mathcal{N}(\mu_3, \Sigma_3); \quad d_\cap^q = d_{\mathcal{N}}(v_\cap, q_\cap) \\ \text{where, } \Sigma_3^{-1} = \Sigma_1^{-1} + \Sigma_2^{-1} \\ \text{and } \mu_3 = \Sigma_3(\Sigma_2^{-1}\mu_1 + \Sigma_1^{-1}\mu_2) \implies \Sigma_3^{-1}\mu_3 = \Sigma_2^{-1}\mu_1 + \Sigma_1^{-1}\mu_2 \end{aligned} \tag{6}$$

Derivation that the intersection of Gaussian density functions is a closed operation[1]. Let us consider two Gaussian PDFs $P(\theta_1) = \mathcal{N}(\mu_1, \Sigma_1)$ and $P(\theta_2) = \mathcal{N}(\mu_2, \Sigma_2)$. Their intersection implies a random variable that is part of both $P(\theta_1)$ and $P(\theta_2)$. The intersection $P(\theta) = \mathcal{N}(\mu_3, \Sigma_3)$ is derived as follows:

$$\begin{aligned} P(\theta) &= P(\theta_1).P(\theta_2) \\ \log(P(\theta)) &= (x-\mu_1)^T\Sigma_1^{-1}(x-\mu_1) + (x-\mu_2)^T\Sigma_2^{-1}(x-\mu_2) \\ (x-\mu_3)^T\Sigma_3^{-1}(x-\mu_3) &= (x-\mu_1)^T\Sigma_1^{-1}(x-\mu_1) + (x-\mu_2)^T\Sigma_2^{-1}(x-\mu_2) \\ \text{Comparing coefficients;} \quad \Sigma_3^{-1} &= \Sigma_1^{-1} + \Sigma_2^{-1}; \quad \mu_3 = \Sigma_3(\Sigma_2^{-1}\mu_1 + \Sigma_1^{-1}\mu_2) \end{aligned}$$

Union ($\cup$). We model the union of multiple entities using Gaussian mixtures. The union of entity density functions given by $e_1, e_2, e_3, ..., e_n$ is defined as $q_\cup$ and distance of resultant entity $v_\cup \in V_\cup$ from the query as $d_\cup^q$ given by:

$$\begin{aligned} q_\cup = \sum_{i=0}^{n} \phi_i \mathcal{N}(\mu_{e_i}, \Sigma_{e_i}); \quad d_\cup^q = \sum_{i=0}^{n} \phi_i d_{\mathcal{N}}(v_\cup, \mathcal{N}(\mu_{e_i}, \Sigma_{e_i})) \\ \text{where, } \phi_i = \frac{exp\left(\mathcal{N}(\mu_{e_i}, \Sigma_{e_i})\right)}{\sum_{j=0}^{n} exp\left(\mathcal{N}(\mu_{e_j}, \Sigma_{e_j})\right)} \end{aligned} \tag{7}$$

$\phi_i \in \Phi$ are the weights for each Gaussian density in the Gaussian mixture, calculated using the self-attention mechanism.

4 Chain Reasoning over Knowledge Graphs

We consider the Gaussian density function (embedding of a single entity) as a special case of Gaussian mixture with a single component. This ensures that all the operations defined in Section 3 are closed under the Gaussian space with an output that is either a single (for translations and intersections) or multi-component Gaussian mixture (for unions). Hence, for chaining the queries, we need to define the logical operators with a Gaussian density and a Gaussian mixture input. In this section, we define the different operators (depicted in Figure 3), in the case of a Gaussian mixture input.

[1]Here, we present a proof sketch. Complete proof is provided in Appendix A

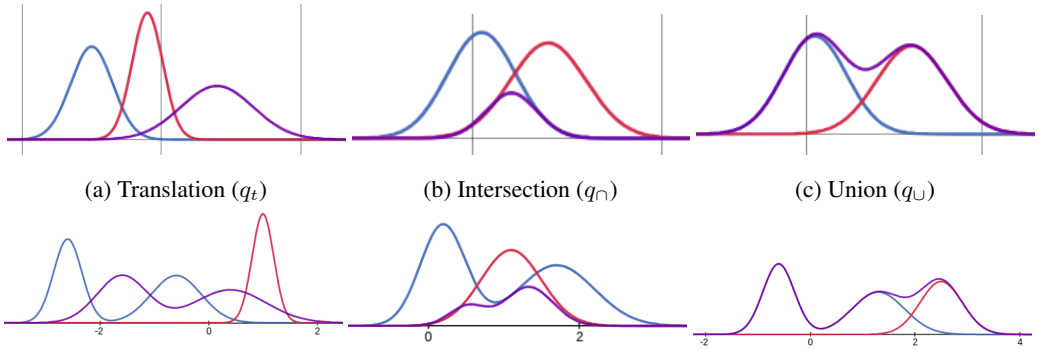

(a) Translation (q_t) (b) Intersection ($q_\cap$) (c) Union ($q_\cup$)

(d) Chain Translation (c_t) (e) Chain Intersection ($c_\cap$) (f) Chain Union ($c_\cup$)

Figure 3: The figures show the logical single and chain operations of translation, intersection, and union in the Gaussian space. The operations are closed and result in either a Gaussian density or a Gaussian mixture. The input operands are given in blue and red and the resultant Gaussian density is depicted in purple. For simplicity, the example is given for a univariate Gaussian model, but in our work, we use multivariate Gaussian densities.

Chain Translation. Let us assume that the input query embedding is an n-component mixture $p = \sum_{i=0}^{n} \mathcal{N}(\mu_i, \Sigma_i)$ and we need to translate it with relation $r = \mathcal{N}(\mu_r, \Sigma_r)$. Intuitively, we would like to translate all the Gaussians in the mixture with the relation. Hence, we model this translation as c_t and the distance from entities $v_t \in V_t$ as d_t^c given by:

$$c_t = \sum_{i=0}^{n} \phi_i \mathcal{N}(\mu_i + \mu_r, (\Sigma_i^{-1} + \Sigma_r^{-1})^{-1}); \;\; d_t^c = \sum_{i=0}^{n} \phi_i d_{\mathcal{N}}(v_t, \mathcal{N}(\mu_i + \mu_r, (\Sigma_i^{-1} + \Sigma_r^{-1})^{-1}) \quad (8)$$

Chain Intersection. A Gaussian mixture is a union over individual densities. Based on the distributive law of sets, an intersection over a Gaussian mixture $p = \sum_{i=0}^{n} \mathcal{N}(\mu_i, \Sigma_i)$ and entity $e = \mathcal{N}(\mu_e, \Sigma_e)$ implies the union of the intersection between the entity and each Gaussian density in the mixture. Hence, we derive this intersection as $c_\cap$ and the distance from entities $v_\cap \in V_\cap$ as $d_\cap^c$:

$$c_\cap = \cup_{i=0}^{n} \mathcal{N}(\mu_e, \Sigma_e)\mathcal{N}(\mu_i, \Sigma_i) = \sum_{i=0}^{n} \phi_i \mathcal{N}(\mu_e, \Sigma_e)\mathcal{N}(\mu_i, \Sigma_i) = \sum_{i=0}^{n} \phi_i \mathcal{N}(\mu_{e\cap i}, \Sigma_{e\cap i}) \quad (9)$$

where, $\Sigma_{e\cap i}^{-1} = \Sigma_e^{-1} + \Sigma_i^{-1}$ and $\mu_{e\cap i} = \Sigma_{e\cap i}(\Sigma_i^{-1}\mu_e + \Sigma_e^{-1}\mu_i)$

$$d_\cap^c = \sum_{i=0}^{n} \phi_i d_{\mathcal{N}}(v_\cap, \mathcal{N}(\mu_{e\cap i}, \Sigma_{e\cap i})) \quad (10)$$

Chain Union. The union of an entity $e = \mathcal{N}(\mu_e, \Sigma_e)$ with a Gaussian mixture $\sum_{i=0}^{n} \phi_i \mathcal{N}(\mu_i, \Sigma_i)$ is the addition of the entity to the mixture. Hence, the union $c_\cup$ and the distance from entities $v_\cup \in V_\cup$ $d_\cup^c$ can be defined as follows:

$$c_\cup = \sum_{i=0}^{n} \phi_i \mathcal{N}(\mu_i, \Sigma_i) + \phi_e \mathcal{N}(\mu_e, \Sigma_e) \quad (11)$$

$$d_\cup^c = \sum_{i=0}^{n} \phi_i d_{\mathcal{N}}(v_\cup, \mathcal{N}(\mu_i, \Sigma_i)) + \phi_e d_{\mathcal{N}}(v_\cup, \mathcal{N}(\mu_e, \Sigma_e)) \quad (12)$$

Implementation Details. To calculate the weights ($\phi_i \in \Phi$) of the Gaussian mixtures, we use the popular self-attention mechanism [21]. The gradient descent over Mahalanobis distance (Eq. 4) and derivation for the product of Gaussians (Eq. 6) are given by [22] and Appendix A, respectively. Another important note is that we do not need the Σ for the operations, but rather Σ^{-1}. Also, storing the complete Σ^{-1} requires quadratic memory, i.e., a Gaussian density of d variables requires $d \times d$ parameters for Σ. So, we only store a decomposed matrix L of Σ^{-1} : $\Sigma^{-1} = LL^T$. Thus, for a Gaussian density of d variables our memory requirement is $2d$ parameters (d for μ and d for Σ^{-1}). For computing the μ_3 for intersection, in Eq. (6), we use a linear solver for faster computation. All our models are implemented in Pytorch [23] and run on two Quadro RTX 8000 [2].

[2] To maintain anonymity, the code is uploaded as supplementary material and a continuously-maintained github link will be provided upon acceptance

5 Experiments

This section describes the experimental setup used to analyze the performance of PERM on various tasks with a focus on the following research questions:

1. Does PERM's query representations perform better than the state-of-the-art baselines on the task of logical reasoning over standard benchmark knowledge graphs?
2. What is the role of individual components in PERM's overall performance gain?
3. Is PERM able to recommend better therapeutic drugs for COVID-19 from a drug re-purposing graph dataset compared to the current baselines?
4. Are we able to visualize the behaviour of PERM's query representations in the latent space?

5.1 Datasets and Baselines

We utilize the following standard benchmark datasets to compare PERM's performance on the task of reasoning over KGs:

- **FB15K-237** [24] is comprised of the 149,689 relation triples and textual mentions of Freebase entity pairs. All the simply invertible relations are removed. The total number of entity and relation types are 14,505 and 474, respectively.
- **NELL995** [25] consists of 107,982 triples obtained from the 995th iteration of the Never-Ending Language Learning (NELL) system. The KG is comprised of 63,361 entities and 400 relations.
- **DBPedia**[3] is a subset of the Wikipedia snapshot that consists of a multi-level hierarchical taxonomy over 240,942 articles. The taxonomy consists of 34,575 entities and 3 relation types.
- **DRKG** [26], or Drug Re-purposing Knowledge Graph (DRKG), is used to evaluate the performance of our model on both the logical reasoning and drug recommendation task. It consists of 5,874,261 relation triples with 97,238 entities and 107 edges.

More detailed statistics of these datasets are provided in Appendix C. For our experiments, we select the following baselines based on (i) their performance on the logical reasoning and (ii) their ability to extend to all FOE query combinations.

- **Graph Query Embedding (GQE)** [6] embeds entities and relations as a vector and utilizes TransE [17] to learn the query embeddings. The distance of the answer entities is calculated using L1-norm.
- **Query2Box (Q2B)** [7] embeds entities and relations as axis aligned hyper-rectangles or boxes and utilize FOE queries to learn query representations. The distance of answer entities is given by a weighted combination of the answer's distance from the center and border of the query box.
- **Beta Query Embedding (BQE)** [11] utilizes beta distribution to learn query representations from FOE queries with a novel addition of negation queries. The distance is calculated as the dimension-wise KL divergence between the answer entity and the query beta embedding.
- **Complex Query Decomposition (CQD)** [10] answers complex queries by reducing them to simpler sub-queries and aggregating the resultant scores with t-norms.

Some of the other baselines [27, 18] focus solely on the multi-hop problem. They could not be intuitively extended to handle all FOE queries, and hence, we did not include them in our study.

5.2 (RQ1) Reasoning over KGs

To evaluate the efficacy of PERM's query representations, we compare it against the baselines on different FOE query types; (i) Single Operator: $1t$, $2t$, $3t$, $2\cap$, $3\cap$, $2\cup$ and (ii) Compound Queries: $\cap t$, $t\cap$, $\cup t$. We follow the standard evaluation protocol [7, 11, 8] and utilize the three splits of a KG for training $\mathcal{G}_{train}$, validation $\mathcal{G}_{valid}$ and evaluation $\mathcal{G}_{test}$ (details in Appendix C). The models are trained on $\mathcal{G}_{train}$ with validation on $\mathcal{G}_{valid}$. The final evaluation metrics for comparison are calculated on $\mathcal{G}_{test}$. For the baselines, we calculate the relevance of the answer entities to the queries based on the distance measures proposed in their respective papers. In PERM, the distance of the

[3] `https://www.kaggle.com/danofer/dbpedia-classes`

answer entity from the query Gaussian density is computed according to the measures discussed in sections 3 and 4. We use the evaluation metrics of HITS@K and MRR to compare the ranked results of the different models. Given the ground truth is $\hat{E}$ and model outputs are $\{e_1, e_2, ..., e_n\} \in E$, the metrics are calculated as:

$$\text{HITS@K} = \frac{1}{K}\sum_{k=1}^{K} f(e_k);\ f(e_k) = \begin{cases} 1, \text{if } e_k \in \hat{E} \\ 0, \text{ else} \end{cases} \quad \text{MRR} = \frac{1}{n}\sum_{i=1}^{n} \frac{1}{f(e_i)};\ f(e_i) = \begin{cases} i, \text{if } e_i \in \hat{E} \\ \infty, \text{ else} \end{cases}$$

Table 1: Performance comparison of PERM (ours) against the baselines to study the efficacy of the query representations. The columns present the different query structures and the overall average performance. The last row presents the Average Relative Improvement (%) of PERM compared to CQD over all datasets across different query types. Best results for each dataset are shown in bold. The MRR results for experiments are given in Appendix D.

Metrics		**HITS@3**									
Dataset	**Model**	**1t**	**2t**	**3t**	**2∩**	**3∩**	**2∪**	**∩*t***	***t*∩**	**∪*t***	**Avg**
FB15k-237	GQE	.404	.214	.147	.262	.390	.164	.087	.162	.155	.221
	BQE	.390	.109	.100	.288	.425	.124	.224	.126	.097	.209
	Q2B	.467	.240	.186	.324	.453	.239	.050	.108	.193	.251
	CQD	.512	**.288**	**.221**	.352	.457	.284	**.129**	**.249**	.121	.290
	PERM	**.520**	.286	.216	**.361**	**.490**	**.305**	.128	.212	**.239**	**.306**
NELL995	GQE	.417	.231	.203	.318	.454	.200	.081	.188	.139	.248
	BQE	.530	.130	.114	**.376**	.475	.122	.129	.241	.086	.246
	Q2B	.555	.266	.233	.343	.480	.369	.132	.212	.163	.306
	CQD	.567	**.294**	**.253**	.350	.476	.441	.131	**.233**	.109	.317
	PERM	**.581**	.286	.243	.352	**.508**	**.460**	**.143**	.195	**.200**	**.328**
DBPedia	GQE	.673	.006	N.A.	.873	.879	.402	.160	.668	0.00	.458
	BQE	.657	.006	N.A.	.964	.966	.306	**.419**	.527	0.00	.481
	Q2B	.832	**.007**	N.A.	**1.00**	**1.00**	.649	.224	.856	0.00	.571
	CQD	.870	**.007**	N.A.	**1.00**	**1.00**	.673	.218	.787	0.00	.569
	PERM	**.950**	**.007**	N.A.	**1.00**	**1.00**	**.782**	.232	**.952**	0.00	**.615**
DRKG	GQE	.420	.218	.153	.270	.409	.181	.101	.186	.174	.235
	BQE	.413	.118	.106	.298	.451	.147	**.270**	.154	.117	.230
	Q2B	.499	.263	.199	.337	.489	.284	.068	.134	.235	.279
	CQD	.554	**.323**	**.238**	.369	.495	.341	.184	**.310**	.150	.329
	PERM	**.565**	.322	.236	**.387**	**.540**	**.376**	.190	.273	**.297**	**.354**
PERM vs Q2B (%)		10.9	12.3	13.0	7.20	6.10	26.3	84.2	50.8	24.3	15.9
PERM vs CQD (%)		3.80	-0.9	-2.4	2.00	5.80	9.50	1.80	-5.5	93.0	6.2

From the results provided in Table 1, we observe that PERM, is able to outperform all the current state-of-the-art approaches, on an average across all FOE queries by 6.2%. Specifically, we see a consistent improvement for union queries; 9.5% and 93% in case of $2\cup$ and $\cup t$, respectively. Comparing the models based on only geometries, we notice the clear efficacy of PERM query representations with an average improvement of 37.9%, 15.9% and 37.3% over vectors (GQE), boxes (Q2B), and beta distribution (BQE), respectively. Given these improvements and the ability to handle compound queries in an end-to-end objective, we conclude that Gaussian distributions are better at learning query representations for FOE reasoning over KGs.

5.3 (RQ2) Ablation Study

In this section, we evaluate the need for different components and their effects on the overall performance of our model. First, we look at the contribution of utilizing different types of queries to the performance of our model. For this, we train our model on different subsets of queries; (i) only $1t$ queries, (ii) only translation ($1t$,$2t$,$3t$) queries and (iii) only single operator queries ($1t$,$2t$,$3t$,$2\cap$,$3\cap$,$2\cup$). Futhermore, we look at the need for attentive aggregation in the case of union of Gaussian mixtures. We test out other methods of aggregation; (i) vanilla averaging and (ii) MLP [28].

From the Table 2, we notice that utilizing only 1t queries significantly reduces the performance of our model by 22.3% and even increasing the scope to all translation queries is still lower in performance

Table 2: Ablation study results. Performance comparison of PERM (final) against different variants of our model. *1t*, *translation* and *single* utilize the 1-hop queries, all translation queries and all single operator queries, respectively. The *average* and *MLP* variants utilize vanilla averaging and MLP for aggregation in union queries. The metrics reported here are an average over all the datasets. Finer evaluation with results for each dataset is given in Appendix E. Best results are given in bold.

Metrics	**HITS@3**									
Model Variants	**1t**	**2t**	**3t**	**2∩**	**3∩**	**2∪**	**∩*t***	***t*∩**	**∪*t***	**Avg**
PERM-1t	.649	.141	.128	.410	.466	.477	.095	.257	.102	.303
PERM-translation	.649	.182	.179	.463	.535	.479	.128	.308	.143	.341
PERM-single	.652	**.225**	.228	.524	.632	.475	.167	.398	.181	.387
PERM-average	.628	.222	.224	.524	.624	.444	.158	.387	.180	.377
PERM-MLP	.642	**.225**	.228	**.526**	.631	.462	.166	.400	.183	.385
PERM (final)	**.654**	**.225**	**.232**	.525	**.635**	**.481**	**.170**	**.408**	**.184**	**.390**

by 12.5% for this case. However, we notice that training on all single operator queries results in comparable performance to the final PERM model. But, given the better overall performance, we utilize all the queries in our final model. For union aggregation, we observe that attention has a clear advantage and both vanilla averaging and MLP lead to a lower performance by 3.33% and 1.28%, respectively. Thus, we adopt self-attention in our final model.

5.4 (RQ3) Case Study: Drug Recommendation

In this experiment, we utilize the expressive power of PERM's query representations to recommend therapeutic drugs for COVID-19 from the DRKG dataset. Drugs in the dataset are already approved for other diseases and the aim is to utilize the drug-protein-disease networks and employ them towards treating COVID-19. This reduces both the drug development time and cost [29]. For this experiment, we utilize the treatment relation in DRKG and retrieve drugs $D : D \xrightarrow{treats} X$, where X is a set of SARS diseases related to the COVID-19 virus. Given that we only need these limited set of entity types (only SARS diseases and drugs) and relation types (only treatments), we only consider the DRKG sub-graph that contains this necessary set of entities and relations for learning the representations. We compare the recommendations of different models against a set of actual candidates currently in trials for COVID-19. We use the top-10 recommendations with the evaluation metrics of precision, recall and F1-score for comparison.

Table 3: The performance comparison of various models on the COVID-19 drug recommendation problem using precision (P), recall (R), and F1-score (F1) metrics. The top three drugs recommended by the models are given in the final column. Recommendations given in green and red indicate a correct and incorrect prediction, respectively. The last two rows provide the Average Relative Improvement of PERM compared to the state-of-the-art baselines Q2B and CQD.

Model	**P@10**	**R@10**	**F1**	**Top Recommended Drugs**
GQE	.119	.174	.141	Piclidenoson, Ibuprofen, Chloroquine
BQE	.159	.200	.177	Ribavirin,Oseltamivir, Ruxolitinib
Q2B	.194	.255	.221	Ribavirin, Dexamethasone, Deferoxamine
CQD	.209	.260	.232	Ribavirin, Dexamethasone, Tofacitinib
PERM	**.217**	**.269**	**.251**	Ribavirin, Dexamethasone, Hydroxychloroquine
PERM vs Q2B (%)	11.9	5.5	13.6	
PERM vs CQD (%)	3.8	3.5	8.2	

We observe from Table 3 that PERM is able to provide the best drug recommendations, across all evaluation metrics. Our model is able to outperform the current methods by atleast 3.8%, 3.5%, and 8.2% in precision, recall, and F1, respectively. Also, the top recommended drugs by our PERM are more inline with the current drug development candidates, thus, showing the better performance of our model's query representations.

5.5 (RQ4) Visualization of the Gaussian Representations

To visualize the entity and query in the latent space, we extract representative entity samples from the FB15K-237 dataset and present them in a 2-dimensional space for better comprehension.

Figure 4 depicts the different entities and the mechanism through which PERM narrows down to the particular answer set. Notice that, we are able to perform an intersection after a union operation due to the closed form nature of our operations. This is currently not possible in state-of-the-art baseline methods. Additionally, it should be noted that, unions widen the query space and intersections narrow them down (as expected). Furthermore, the variance parameter acts as a control over the spatial area that an entity should cover and more general entities such as `Turing Award` and `Europe` occupy a larger area than their respective sub-categories, namely, `winners` and `Europeans`.

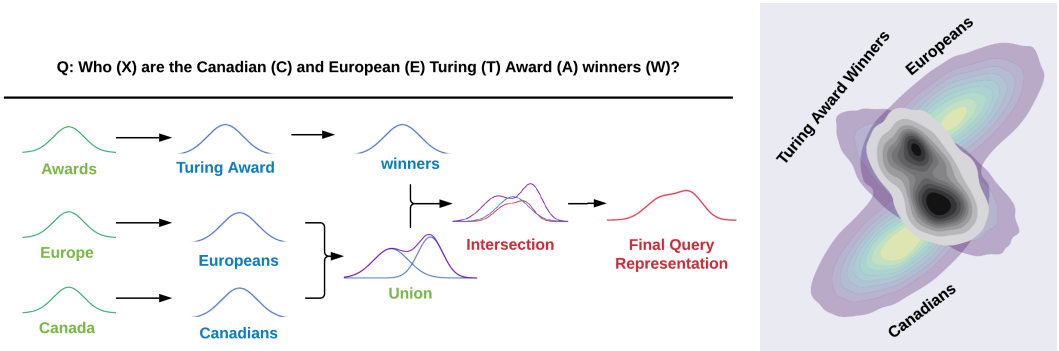

(a) Query processing in PERM. This figure depicts a univariate version of the entity Gaussian embeddings for better visualization of the process. The same property, however, generalizes over an increased number of dimensions, i.e., multivariate case.

(b) Bivariate version of the final query space, given in grayscale with darker colors representing a higher probability of answers.

Figure 4: An illustration of the flow for a sample complex query in the representational space. We note that intersection after union is possible in our PERM model because the operations are closed and this is not possible in current methods including BQE, Q2B, and CQD.

6 Conclusion

In this paper, we advocate for Probabilistic Entity Representation Model (PERM) to learn query representations for chain reasoning over knowledge graphs. We show the representational power of our model by defining closed solutions to FOE queries and their chains. Additionally, we also demonstrate its performance compared to its state-of-the-art counterparts on the problems of reasoning over KGs and drug recommendation for COVID-19 from the DRKG dataset. Furthermore, we exhibit its interpretability by depicting the representational space through a sample query processing pipeline.

7 Broader Impact

PERM is the first method that models an individual entity in knowledge graphs using Gaussian density function, making it possible to solve FOE queries in a closed form solution. This enables its application in domains that require chain reasoning. The basic idea behind the solution can also be extended to any domain that can encode its basic units as Gaussians and extend the units through FOE queries, e.g., in topic modeling, topics can be encoded as Gaussians and documents as union of topics.

However, PERM depends on the integrity of the knowledge graph used for training. Any malicious attacks/errors [30, 31] that lead to incorrect relations could, further, lead to incorrect results and affect the confidence of our model. Furthermore, due to the connected nature of complex queries, this attack could propagate and affect a larger set of queries. Such incorrect results would be fatal in sensitive areas of research such as drug recommendations and, thus, it is necessary to maintain the integrity of training data before learning representations and querying with PERM.

References

[1] Linfeng Li, Peng Wang, Jun Yan, Yao Wang, Simin Li, Jinpeng Jiang, Zhe Sun, Buzhou Tang, Tsung-Hui Chang, Shenghui Wang, and Yuting Liu. Real-world data medical knowledge graph: construction and applications. *Artificial Intelligence in Medicine*, 103:101817, 2020.

[2] Longxiang Shi, Shijian Li, Xiaoran Yang, Jiaheng Qi, Gang Pan, and Binbin Zhou. Semantic health knowledge graph: semantic integration of heterogeneous medical knowledge and services. *BioMed research international*, 2017.

[3] Xin Luna Dong, Xiang He, Andrey Kan, Xian Li, Yan Liang, Jun Ma, Yifan Ethan Xu, Chenwei Zhang, Tong Zhao, Gabriel Blanco Saldana, et al. Autoknow: Self-driving knowledge collection for products of thousands of types. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2724–2734, 2020.

[4] Christian Bizer, Jens Lehmann, Georgi Kobilarov, Sören Auer, Christian Becker, Richard Cyganiak, and Sebastian Hellmann. Dbpedia-a crystallization point for the web of data. *Journal of web semantics*, 7(3):154–165, 2009.

[5] Piero Andrea Bonatti, Stefan Decker, Axel Polleres, and Valentina Presutti. Knowledge Graphs: New Directions for Knowledge Representation on the Semantic Web (Dagstuhl Seminar 18371). *Dagstuhl Reports*, 8(9):29–111, 2019.

[6] William L. Hamilton, Payal Bajaj, Marinka Zitnik, Dan Jurafsky, and Jure Leskovec. Embedding logical queries on knowledge graphs. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, page 2030–2041, Red Hook, NY, USA, 2018. Curran Associates Inc.

[7] Hongyu Ren*, Weihua Hu*, and Jure Leskovec. Query2box: Reasoning over knowledge graphs in vector space using box embeddings. In *International Conference on Learning Representations*, 2020.

[8] Nurendra Choudhary, Nikhil Rao, Sumeet Katariya, Karthik Subbian, and Chandan K Reddy. Self-supervised hyperboloid representations from logical queries over knowledge graphs. *arXiv preprint arXiv:2012.13023*, 2020.

[9] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. Complex embeddings for simple link prediction. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML'16, page 2071–2080. JMLR.org, 2016.

[10] Erik Arakelyan, Daniel Daza, Pasquale Minervini, and Michael Cochez. Complex query answering with neural link predictors. In *International Conference on Learning Representations*, 2021.

[11] Hongyu Ren and Jure Leskovec. Beta embeddings for multi-hop logical reasoning in knowledge graphs. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 19716–19726. Curran Associates, Inc., 2020.

[12] Katrin Erk. Representing words as regions in vector space. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL-2009)*, pages 57–65, 2009.

[13] Luke Vilnis, Xiang Li, Shikhar Murty, and Andrew McCallum. Probabilistic embedding of knowledge graphs with box lattice measures. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 263–272, Melbourne, Australia, July 2018. Association for Computational Linguistics.

[14] Luke Vilnis and Andrew McCallum. Word representations via gaussian embedding. In *ICLR*, 2015.

[15] Aleksandar Bojchevski and Stephan Günnemann. Deep gaussian embedding of graphs: Unsupervised inductive learning via ranking. In *International Conference on Learning Representations*, 2018.

[16] Yadolah Dodge. *Mahalanobis Distance*, pages 325–326. Springer New York, New York, NY, 2008.

[17] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. In *Neural Information Processing Systems (NIPS)*, pages 1–9, 2013.

[18] Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. A three-way model for collective learning on multi-relational data. In *Icml*, 2011.

[19] Rajarshi Das, Arvind Neelakantan, David Belanger, and Andrew McCallum. Chains of reasoning over entities, relations, and text using recurrent neural networks. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 132–141, Valencia, Spain, April 2017. Association for Computational Linguistics.

[20] Haitian Sun, Andrew O Arnold, Tania Bedrax-Weiss, Fernando Pereira, and William W Cohen. Faithful embeddings for knowledge base queries. *Advances in Neural Information Processing Systems*, 33, 2020.

[21] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

[22] Marc T Law, Yaoliang Yu, Matthieu Cord, and Eric P Xing. Closed-form training of mahalanobis distance for supervised clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3909–3917, 2016.

[23] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.

[24] Kristina Toutanova, Danqi Chen, Patrick Pantel, Hoifung Poon, Pallavi Choudhury, and Michael Gamon. Representing text for joint embedding of text and knowledge bases. In *Proceedings of the 2015 conference on empirical methods in natural language processing*, pages 1499–1509, 2015.

[25] Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam Hruschka, and Tom Mitchell. Toward an architecture for never-ending language learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 24, 2010.

[26] Vassilis N. Ioannidis, Xiang Song, Saurav Manchanda, Mufei Li, Xiaoqin Pan, Da Zheng, Xia Ning, Xiangxiang Zeng, and George Karypis. Drkg - drug repurposing knowledge graph for covid-19. `https://github.com/gnn4dr/DRKG/`, 2020.

[27] Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. *arXiv preprint arXiv:1412.6575*, 2014.

[28] Fionn Murtagh. Multilayer perceptrons for classification and regression. *Neurocomputing*, 2(5):183–197, 1991.

[29] Sudeep Pushpakom, Francesco Iorio, Patrick A Eyers, K Jane Escott, Shirley Hopper, Andrew Wells, Andrew Doig, Tim Guilliams, Joanna Latimer, Christine McNamee, et al. Drug repurposing: progress, challenges and recommendations. *Nature reviews Drug discovery*, 18(1):41–58, 2019.

[30] Daniel Zügner, Amir Akbarnejad, and Stephan Günnemann. Adversarial attacks on neural networks for graph data. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2847–2856, 2018.

[31] Hanjun Dai, Hui Li, Tian Tian, Xin Huang, Lin Wang, Jun Zhu, and Le Song. Adversarial attack on graph structured data. In *International conference on machine learning*, pages 1115–1124. PMLR, 2018.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Section xx.
- Did you include the license to the code and datasets? [No] The code and the data are proprietary.
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] Refer section 5.
   (b) Did you describe the limitations of your work? [Yes] Refer section 7.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] Refer section 7.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [Yes] Refer Appendix A.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] As supplementary material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Refer Appendix C for data splits and section 4 for implementation details of hyperparameters and computational resources.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] The experiments are computationally expensive and all the baselines and our model have been run multiple times for best results.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] Refer to section 4 for implementation details.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] All existing assets have been cited in the relevant places.
   (b) Did you mention the license of the assets? [Yes] Given with the code in the supplementary material. All existing assets are public open-source libraries.
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] Included in the supplementary material.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]