
A Contextualized BERT model for Knowledge Graph Completion

Anonymous Author(s)

Abstract

1 Knowledge graphs (KGs) are valuable for representing structured, interconnected
2 information across domains, enabling tasks like semantic search, recommendation
3 systems and inference. A pertinent challenge with KGs, however, is that many
4 entities (i.e., heads, tails) or relationships are unknown. Knowledge Graph Com-
5 pletion (KGC) addresses this by predicting these missing nodes or links, enhancing
6 the graph’s informational depth and utility. Traditional methods like TransE and
7 ComplEx predict tail entities but struggle with unseen entities. Textual-based
8 models leverage additional semantics but come with high computational costs,
9 semantic inconsistencies, and data imbalance issues. Recent LLM-based models
10 show improvement but overlook contextual information and rely heavily on entity
11 descriptions. In this study, we introduce a contextualized BERT model for KGC
12 that overcomes these limitations by utilizing the contextual information from neigh-
13 bouring entities and relationships to predict tail entities. Our model eliminates the
14 need for entity descriptions and negative triplet sampling, reducing computational
15 demands while improving performance. Our model outperforms state-of-the-art
16 methods on standard datasets, improving Hit@1 by 5.3% and 4.88% on FB15k-237
17 and WN18RR respectively, setting a new benchmark in KGC.¹

18 1 Introduction

19 A knowledge graph (KG) is a structured representation of entities (as nodes) and relationships (as
20 links) that supports search, recommendation and other downstream reasoning tasks. However, KGs
21 are often incomplete, with many entities (heads/tails) or relationships missing, limiting their utility
22 in real-world applications [1]. Consequently, Knowledge Graph Completion (KGC)—predicting a
23 missing tail entity $(h, r, ?)$, head entity $(?, r, t)$, or relationship $(h, ?, t)$ in a triplet—has become a
24 critical research objective, with numerous methodologies proposed to tackle this issue.

25 Embedding-based methods, for instance, learn vector embeddings for entities and relationships from
26 training data, but these methods struggle to generalize to *unseen* entities or relationships, impairing
27 performance in tail prediction during testing [2]. Recently, large language model (LLM)-based
28 approaches for KGC have shown potential in overcoming this limitation by leveraging LLMs trained
29 on extensive datasets to capture complex semantic relationships and generalize better to unseen
30 entities [3, 4, 5]. Despite these strengths, LLM-based models are computationally demanding, often
31 overlook relational context and depend heavily on entity descriptions and negative sampling. More
32 recent LLM prompting-based approaches encode KGs into prompts[4], but injecting all relevant
33 facts from a KG into prompts is labor-intensive, and generic LLMs often struggle with domain-
34 specific KGs. Additionally, textual information-based methods like NN-KGC and Sim-KGC utilize
35 neighborhood information for KGC, but they often require entity descriptions, which may not be
36 available in many datasets, and add computational overhead [6, 7]. To address these limitations, we
37 propose a Context-Aware BERT for Knowledge Graph Completion (CAB-KGC) [9] that extracts

¹The main paper for this work has been submitted to ICLR 2025 for consideration.

38 contextual information associated with the operational relationship, its neighboring entities and
 39 relationships associated with the head entity. This context is then integrated with the BERT model to
 40 enhance the prediction of tail entities. To summarise, this study makes the following contributions to
 41 the KG domain:

- 42 • We introduce the CAB-KGC approach [9] to address the KGC problem, leveraging graph
 43 features of head entity context and relationship context, and BERT model. The CAB-KGC
 44 approach outperforms SOTA KGC methods.
- 45 • CAB-KGC eliminates reliance on entity descriptions, focusing solely on head and relation-
 46 ship contexts for improved predictions available in all KGs.
- 47 • CAB-KGC does not require negative sample training, enhancing training speed and resilience
 48 against negative sample selection.
- 49 • Extensive experiments on various benchmark datasets demonstrate that CAB-KGC reliably
 50 excels in tail entity prediction.

51 2 Methodology

52 Problem Formulation (see Table 1 for notations): Consider a knowledge graph $G(E, R)$ as a collection
 53 of triplets (h, r, t) , where $h \in E$ is the head entity, $t \in E$ is the tail entity, and $r \in R$ represents the
 54 relationship between them, our CAB-KGC model predicts a missing tail t (represented by ?) given
 55 an incomplete triple $(h, r, ?)$.

Table 1: Mathematical Notations and Symbols

Notation	Description	Notation	Description
e	entity or node	r	relationship
h	head entity node	t	tail entity node
E	Entities Set	R	Relationships Set
H_c	Head (h) or Entity context	R_c	Relationship context
$R(h), E(h)$	relation and entities associated to head h	N_T	Total number of triplets
$p_{\theta}(t_i h_i, r_i)$	Tail t_i probability given head h_i and relationship r_i	rank_i	Rank of the true tail entity t_i in the prediction
↓	Results: Lower is better	↑	Results: Higher is better

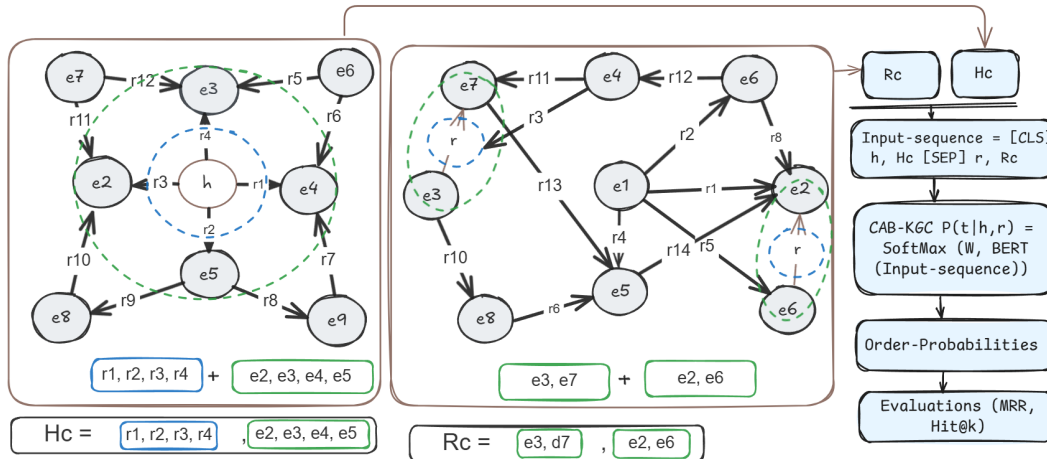


Figure 1: A concise view of the CAB-KGC Method. Box on the left shows head context H_c calculation, the middle one shows relationship context R_c calculation. H_c and R_c are then fed into the model pipeline shown on the right side.

56 Figure 1 provides an overview of the CAB-KGC model. It predicts the tail entity t given a head h
 57 and a relationship r , in the following steps:

- 58 1. **Extract Head Context H_c :** To extract the contextual information for the head i.e. H_c , we
 59 first identify the relationships r that are associated with the head entity h , i.e., $R(h)$. If k
 60 relationships are associated with the head h from the set R of all relationships r_i in the
 61 graph G , then:

$$R(h) = \text{set}\{r_1, r_2, \dots, r_k\} \quad (1)$$

62 Next, we find the entities e that are neighbours (have a direct connection) with the head
 63 entity h , i.e., $E(h)$ using the identified relationships $R(h)$. These neighbour entities can be
 64 mathematically expressed as:

$$E(h) = \text{set}\{e_1, e_2, \dots, e_m\} \quad (2)$$

65 The head context H_c is then calculated as the union of the connected relationships $R(h)$ and
 66 the neighbour entities $E(h)$, as shown below in Equation 3.

$$H_c = (R(h)) \cup (E(h)) \quad (3)$$

68 2. **Extract Relationship Context R_c :** To acquire the relationship context R_c , we identify all
 69 the entities associated with the operational relationship r in the knowledge Graph G . R_c is
 70 given as:

$$R_c = \text{set}\{e_1, e_2, \dots, e_l\} \quad (4)$$

71 3. **Prepare Input Sequence for BERT Classifier:** The contextual information extracted in
 72 above steps forms the input to BERT. Specifically, the input sequence contains h , H_c from
 73 Equation 3, r , and R_c from Equation 4, as shown below:

$$\text{Input Sequence} = [CLS] h, H_c [SEP] r, R_c \quad (5)$$

74 where [CLS] is BERT’s classifier token and [SEP] is the separator token.

75 4. **Predict and train with BERT Classifier:** A classification layer is added on top of the BERT
 76 model, which aims to classify the tail entity ($h, r, ?$). Once the BERT classifier receives
 77 the input, it processes it through various transformer layers, provides a contextualized
 78 representation of each token and uses that to classify the input. The classifier model predicts
 79 the tail entity by employing a softmax function over the output embedding to calculate the
 80 probability for all the available tail entities. The input-output description of the model is
 81 given as:

$$(\text{BERT Output}) = \text{BERT}(\text{Input Sequence}) \quad (6)$$

$$P(t | h, r) = \text{softmax}(W \cdot \text{BERT Output}) \quad (7)$$

83 where W is a learned weight matrix.

84 Putting above equations together, the CAB-KGC model can be expressed as:

$$\text{CAB-KGC}(t | h, r) = \text{softmax}(W \cdot \text{BERT}(h, H_c, r, R_c)) \quad (8)$$

85 The CAB-KGC model is trained using cross-entropy loss, which compares the probability
 86 distribution of the predicted label with the true label for the tail entity. The cross-entropy
 87 loss is given by:

$$L = - \sum_{i=1}^N y_i \log P(t_i | h, r) \quad (9)$$

88 In this equation the one-hot encoded true label for the tail object t_i is indicated as y_i . The
 89 predicted probability for the true tail entity could be denoted as $P(t_i | h, r)$, where h is the
 90 head and r is the relation.

91 2.1 Experiments Setup

92 **Datasets:** We assessed the proposed CAB-KGC model on various commonly used KG datasets.
 93 These datasets are briefly explained here:

- 94 • FB15k-237 [22] is an updated version subset of the FB15k dataset, where the inverse triplets
 95 have been removed to increase the difficulty of the KGC. It has 14541 unique entities and
 96 237 relationships.
- 97 • WN18RR [23] is the subset of WN18, where the reverse triplets are removed, making it
 98 more complex for the models to incorporate the problem of KGC.

99 **Hyperparameters:** The experiments used a batch size of 16 and a learning rate of 5e-5, Adam as the
 100 optimizer and cross-entropy as the loss function. The experiments were accomplished on an NVIDIA
 101 GeForce RTX 3090 GPU with 24 GB of memory. Training for the CAB-KGC model was halted once
 102 evaluation metrics stabilized to the third decimal place.

103 **Evaluation:** Various standard evaluation metrics in KGC, as given in Equation 10, such as MRR,
 104 and Hit@k, are utilized to assess the performance of the proposed method and other state-of-the-art
 105 approaches

$$\text{MRR} = \frac{1}{N} \sum_{i=1}^N \frac{1}{\text{rank}_i} \quad ; \quad \text{Hits@k} = \frac{1}{N} \sum_{i=1}^N \mathbf{1}(\text{rank}_i \leq k) \quad (10)$$

106 where rank_i is the correct entity rank position in the descending order sorted list of predicted scores
 107 for the i -th triplet. The function $\mathbf{1}(\text{rank}_i \leq k)$ is a ranking function that outputs one if the true entity
 108 is ranked within the top k predictions and 0 otherwise.

109 2.2 Results

110 Our CAB-KGC approach shows superior results on the FB15k-237 dataset. CAB-KGC’s significant
 111 performance is its Hits@1 score of 0.322, which improves SOTA by almost 5.3%, showing a superior
 112 ability to rank accurate entities in the first place. It obtains a Hits@3 score of 0.399 and improves
 113 by 0.5%, notably above other models, indicating that CAB-KGC reliably predicts relevant entities
 114 within the top 3 ranks. The CAB-KGC method performed well on the WN18RR dataset, getting
 115 an MRR of 0.685, which is an improvement of 1.2% over SOTA models and a Hits@1 of 0.637,
 116 an improvement of 4.88%, outperforming state-of-the-art methods. All the results are reported in
 117 Table 2. Note that we have excluded results from KICGPT[4] during comparison for reasons: (a) the
 118 model is prompt-based and not trainable, (b) its performance is highly dependent on the underlying
 119 LLM’s knowledge base and (c) large KGs cannot be injected as prompts to this model. Furthermore,
 120 injecting all relevant facts from different KGs into prompts is labor-intensive, and underlying LLM
 121 will often struggle with domain-specific KGs when it does not contain enough relevant knowledge.

Table 2: Comparison of the proposed and baseline methods on the datasets FB15k-237 and WN18RR. The optimal outcome for each metric is highlighted in bold, while the second-best result is underlined. The circle symbol \circ denotes that the results have been extracted from the study conducted by Wei et al. [8], while the symbol \square indicates that the results have been extracted from the study conducted by Yao et al. in [3].

Dataset	FB15k-237			WN18RR		
	MRR \uparrow	Hits@1 \uparrow	Hits@3 \uparrow	MRR \uparrow	Hits@1 \uparrow	Hits@3 \uparrow
<i>Embedding-Based Methods</i>						
RESCAL [10] \circ	<u>0.356</u>	0.266	0.390	0.467	0.439	0.478
TransE [11] \circ	0.279	0.198	0.376	0.243	0.043	0.441
DistMult [12] \circ	0.241	0.155	0.263	0.430	0.390	0.440
ComplEx [13] \circ	0.247	0.158	0.275	0.440	0.410	0.460
RotatE [14] \circ	0.338	0.241	0.375	0.476	0.428	0.492
TuckER [15] \circ	0.358	<u>0.266</u>	0.394	0.470	0.443	0.482
CompGCN [16] \circ	0.355	0.264	<u>0.390</u>	0.479	0.443	0.494
HittER [17] \circ	0.344	0.246	0.380	0.496	0.449	0.514
HAKE [18] \circ	0.346	0.250	0.381	0.497	0.452	0.516
<i>Text-and Description-Based Methods</i>						
Pretrain-KGE [5] \circ	0.332	-	-	0.235	-	-
StAR [19] \circ	0.263	0.171	0.287	0.364	0.222	0.436
MEM-KGC (w/o EP) [20] \circ	0.339	0.249	0.372	0.533	0.473	0.570
MEM-KGC (w/ EP) [20] \circ	0.346	0.253	0.381	0.557	0.475	0.604
SimKGC [6] \circ	0.333	0.246	0.363	0.671	0.585	0.731
NNKGC [21] \circ	0.338	0.252	0.365	<u>0.674</u>	<u>0.596</u>	<u>0.722</u>
<i>LLM-Based Methods</i>						
ChatGPTzero-shot [5] \square	-	0.237	-	-	0.190	-
ChatGPTone-shot [5] \square	-	0.267	-	-	0.212	-
KICGPT [5] \square	0.410	0.321	0.430	0.564	0.478	0.612
<i>Proposed</i>						
Proposed CAB-KGC	0.350	0.322	0.399	0.685	0.637	0.687

122 3 Conclusion

123 The proposed CAB-KGC approach exploits the contextual information to outperform existing methods
 124 in MRR and Hit@k measures, with improvements of 5.3% and 4.88% over FB15k-237 and WN18RR,
 125 respectively.

References

- 126
- 127 [1] Ding, J. and Jia, W. (2018) The research advances of knowledge graph completion algorithm.
128 *Information and Communication Technology* 1:56–62.
- 129 [2] Xie, R., Liu, Z., Jia, J., Luan, H., and Sun, M. (2016) Representation learning of knowl-
130 edge graphs with entity descriptions. In *Proceedings of the AAAI Conference on Artificial*
131 *Intelligence*, Vol. 30, No. 1.
- 132 [3] Yao, L., Mao, C., and Luo, Y. (2019) KG-BERT: BERT for Knowledge Graph Completion.
133 *arXiv preprint* arXiv:1909.03193. Available at: <https://arxiv.org/abs/1909.03193>.
- 134 [4] Wei, Y., Huang, Q., Zhang, Y., and Kwok, J. (2023) KICGPT: Large Language Model with
135 Knowledge in Context for Knowledge Graph Completion. In *Findings of the Association for*
136 *Computational Linguistics: EMNLP 2023*, pp. 8667–8683. Association for Computational
137 Linguistics. Available at: <http://dx.doi.org/10.18653/v1/2023.findings-emnlp.580>.
- 138 [5] Zhang, Z., Liu, X., Zhang, Y., Su, Q., Sun, X., and He, B. (2020) Pretrain-KGE:
139 Learning Knowledge Representation from Pretrained Language Models. In T. Cohn,
140 Y. He, and Y. Liu (eds.), *Findings of the Association for Computational Linguistics:*
141 *EMNLP 2020*, pp. 259–266. Association for Computational Linguistics. Available at:
142 <https://aclanthology.org/2020.findings-emnlp.25>.
- 143 [6] Wang, L., Zhao, W., Wei, Z., and Liu, J. (2022) SimKGC: Simple Contrastive Knowledge
144 Graph Completion with Pre-trained Language Models. *arXiv preprint* arXiv:2203.02167.
145 Available at: <https://arxiv.org/abs/2203.02167>.
- 146 [7] Li, D., Tan, Z., Chen, T., and Liu, H. (2024) Contextualization Distillation from Large Lan-
147 guage Model for Knowledge Graph Completion. *arXiv preprint* arXiv:2402.01729. Available
148 at: <https://arxiv.org/abs/2402.01729>.
- 149 [8] Wei, Y., Huang, Q., Kwok, J.T., and Zhang, Y. (2024) Kicgpt: Large language model with
150 knowledge in context for knowledge graph completion. *arXiv preprint* arXiv:2402.02389.
- 151 [9] Anonymous Authors. (2025) CAB-KGC: Context-Aware BERT for Knowledge Graph Com-
152 pletion. *Submitted to ICLR 2025 for consideration*
- 153 [10] Nickel, M., Tresp, V., and Kriegel, H.-P. (2011) A three-way model for collective learning
154 on multi-relational data. In *Proceedings of the 28th International Conference on Machine*
155 *Learning*, pp. 809–816. Omnipress. Madison, WI, USA.
- 156 [11] Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., and Yakhnenko, O. (2013) Translating
157 embeddings for modeling multi-relational data. In C.J. Burges, L. Bottou, M. Welling, Z.
158 Ghahramani, and K.Q. Weinberger (eds.), *Advances in Neural Information Processing Systems*,
159 Vol. 26. Curran Associates, Inc.
- 160 [12] Yang, B., Yih, W.-t., He, X., Gao, J., and Deng, L. (2014) Embedding entities and relations
161 for learning and inference in knowledge bases. In *International Conference on Learning*
162 *Representations*. Available at: <https://api.semanticscholar.org/CorpusID:2768038>.
- 163 [13] Trouillon, T., Welbl, J., Riedel, S., Gaussier, É., and Bouchard, G. (2016) Complex
164 embeddings for simple link prediction. *arXiv preprint* arXiv:1606.06357. Available at:
165 <https://arxiv.org/abs/1606.06357>.
- 166 [14] Sun, Z., Deng, Z.-H., Nie, J.-Y., and Tang, J. (2019) RotatE: Knowledge graph embed-
167 ding by relational rotation in complex space. *arXiv preprint* arXiv:1902.10197. Available at:
168 <https://arxiv.org/abs/1902.10197>.
- 169 [15] Wang, Y., Broscheit, S., and Gemulla, R. (2019) A relational Tucker decomposi-
170 tion for multi-relational link prediction. *arXiv preprint* arXiv:1902.00898. Available at:
171 <https://arxiv.org/abs/1902.00898>.
- 172 [16] Vashishth, S., Sanyal, S., Nitin, V., and Talukdar, P. (2020) Composition-based multi-
173 relational graph convolutional networks. *arXiv preprint* arXiv:1911.03082. Available at:
174 <https://arxiv.org/abs/1911.03082>.

- 175 [17] Chen, S., Liu, X., Gao, J., Jiao, J., Zhang, R., and Ji, Y. (2020) Hitter: Hierarchical transformers
176 for knowledge graph embeddings. *arXiv preprint* arXiv:2008.12813.
- 177 [18] Zhang, Z., Cai, J., Zhang, Y., and Wang, J. (2022) Learning hierarchy-aware knowl-
178 edge graph embeddings for link prediction. *arXiv preprint* arXiv:1911.09419. Available at:
179 <https://arxiv.org/abs/1911.09419>.
- 180 [19] Wang, B., Shen, T., Long, G., Zhou, T., Wang, Y., and Chang, Y. (2021) Structure-augmented
181 text representation learning for efficient knowledge graph completion. In *Proceedings of the*
182 *Web Conference 2021*, pp. 1737–1748.
- 183 [20] Choi, B., Jang, D., and Ko, Y. (2021) MEM-KGC: Masked entity model for knowledge graph
184 completion with pre-trained language model. *IEEE Access* **9**, 132025-132032. Available at:
185 <https://api.semanticscholar.org/CorpusID:238243180>.
- 186 [21] Li, I., and Yang, B. (2023) NNKGC: Improving knowledge graph completion with node neigh-
187 borhoods. *arXiv preprint* arXiv:2302.06132. Available at: <https://arxiv.org/abs/2302.06132>.
- 188 [22] Bollacker, K., Evans, C., Paritosh, P., Sturge, T., and Taylor, J. (2008) Freebase: a collabora-
189 tively created graph database for structuring human knowledge. In *Proceedings of the 2008*
190 *ACM SIGMOD International Conference on Management of Data*, pp. 1247–1250.
- 191 [23] Miller, G. A. (1995) WordNet: a lexical database for English. *Communications of the ACM*
192 **38**(11):39–41.

193 **NeurIPS Paper Checklist**

194 **1. Claims**

195 Question: Do the main claims made in the abstract and introduction accurately reflect the
196 paper's contributions and scope?

197 Answer: [Yes]

198 Guidelines:

- 199 • The answer NA means that the abstract and introduction do not include the claims
200 made in the paper.
- 201 • The abstract and/or introduction should clearly state the claims made, including the
202 contributions made in the paper and important assumptions and limitations. A No or
203 NA answer to this question will not be perceived well by the reviewers.
- 204 • The claims made should match theoretical and experimental results, and reflect how
205 much the results can be expected to generalize to other settings.
- 206 • It is fine to include aspirational goals as motivation as long as it is clear that these goals
207 are not attained by the paper.

208 **2. Limitations**

209 Question: Does the paper discuss the limitations of the work performed by the authors?

210 Answer: [Yes]

211 Guidelines:

- 212 • The answer NA means that the paper has no limitation while the answer No means that
213 the paper has limitations, but those are not discussed in the paper.
- 214 • The authors are encouraged to create a separate "Limitations" section in their paper.
- 215 • The paper should point out any strong assumptions and how robust the results are to
216 violations of these assumptions (e.g., independence assumptions, noiseless settings,
217 model well-specification, asymptotic approximations only holding locally). The authors
218 should reflect on how these assumptions might be violated in practice and what the
219 implications would be.
- 220 • The authors should reflect on the scope of the claims made, e.g., if the approach was
221 only tested on a few datasets or with a few runs. In general, empirical results often
222 depend on implicit assumptions, which should be articulated.
- 223 • The authors should reflect on the factors that influence the performance of the approach.
224 For example, a facial recognition algorithm may perform poorly when image resolution
225 is low or images are taken in low lighting. Or a speech-to-text system might not be
226 used reliably to provide closed captions for online lectures because it fails to handle
227 technical jargon.
- 228 • The authors should discuss the computational efficiency of the proposed algorithms
229 and how they scale with dataset size.
- 230 • If applicable, the authors should discuss possible limitations of their approach to
231 address problems of privacy and fairness.
- 232 • While the authors might fear that complete honesty about limitations might be used by
233 reviewers as grounds for rejection, a worse outcome might be that reviewers discover
234 limitations that aren't acknowledged in the paper. The authors should use their best
235 judgment and recognize that individual actions in favor of transparency play an impor-
236 tant role in developing norms that preserve the integrity of the community. Reviewers
237 will be specifically instructed to not penalize honesty concerning limitations.

238 **3. Theory Assumptions and Proofs**

239 Question: For each theoretical result, does the paper provide the full set of assumptions and
240 a complete (and correct) proof?

241 Answer: [Yes]

242 Guidelines:

- 243 • The answer NA means that the paper does not include theoretical results.
- 244 • All the theorems, formulas, and proofs in the paper should be numbered and cross-
245 referenced.

- 246 • All assumptions should be clearly stated or referenced in the statement of any theorems.
- 247 • The proofs can either appear in the main paper or the supplemental material, but if
- 248 they appear in the supplemental material, the authors are encouraged to provide a short
- 249 proof sketch to provide intuition.
- 250 • Inversely, any informal proof provided in the core of the paper should be complemented
- 251 by formal proofs provided in appendix or supplemental material.
- 252 • Theorems and Lemmas that the proof relies upon should be properly referenced.

253 4. Experimental Result Reproducibility

254 Question: Does the paper fully disclose all the information needed to reproduce the main ex-
255 perimental results of the paper to the extent that it affects the main claims and/or conclusions
256 of the paper (regardless of whether the code and data are provided or not)?

257 Answer: [Yes]

258 Guidelines:

- 259 • The answer NA means that the paper does not include experiments.
- 260 • If the paper includes experiments, a No answer to this question will not be perceived
- 261 well by the reviewers: Making the paper reproducible is important, regardless of
- 262 whether the code and data are provided or not.
- 263 • If the contribution is a dataset and/or model, the authors should describe the steps taken
- 264 to make their results reproducible or verifiable.
- 265 • Depending on the contribution, reproducibility can be accomplished in various ways.
- 266 For example, if the contribution is a novel architecture, describing the architecture fully
- 267 might suffice, or if the contribution is a specific model and empirical evaluation, it may
- 268 be necessary to either make it possible for others to replicate the model with the same
- 269 dataset, or provide access to the model. In general, releasing code and data is often
- 270 one good way to accomplish this, but reproducibility can also be provided via detailed
- 271 instructions for how to replicate the results, access to a hosted model (e.g., in the case
- 272 of a large language model), releasing of a model checkpoint, or other means that are
- 273 appropriate to the research performed.
- 274 • While NeurIPS does not require releasing code, the conference does require all submis-
275 sions to provide some reasonable avenue for reproducibility, which may depend on the
276 nature of the contribution. For example
- 277 (a) If the contribution is primarily a new algorithm, the paper should make it clear how
278 to reproduce that algorithm.
- 279 (b) If the contribution is primarily a new model architecture, the paper should describe
280 the architecture clearly and fully.
- 281 (c) If the contribution is a new model (e.g., a large language model), then there should
282 either be a way to access this model for reproducing the results or a way to reproduce
283 the model (e.g., with an open-source dataset or instructions for how to construct
284 the dataset).
- 285 (d) We recognize that reproducibility may be tricky in some cases, in which case
286 authors are welcome to describe the particular way they provide for reproducibility.
287 In the case of closed-source models, it may be that access to the model is limited in
288 some way (e.g., to registered users), but it should be possible for other researchers
289 to have some path to reproducing or verifying the results.

290 5. Open access to data and code

291 Question: Does the paper provide open access to the data and code, with sufficient instruc-
292 tions to faithfully reproduce the main experimental results, as described in supplemental
293 material?

294 Answer: [Yes]

295 Guidelines:

- 296 • The answer NA means that paper does not include experiments requiring code.
- 297 • Please see the NeurIPS code and data submission guidelines ([https://nips.cc/
298 public/guides/CodeSubmissionPolicy](https://nips.cc/public/guides/CodeSubmissionPolicy)) for more details.
- 299 • While we encourage the release of code and data, we understand that this might not be
300 possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not

- 301 including code, unless this is central to the contribution (e.g., for a new open-source
302 benchmark).
- 303 • The instructions should contain the exact command and environment needed to run to
304 reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
 - 305 • The authors should provide instructions on data access and preparation, including how
306 to access the raw data, preprocessed data, intermediate data, and generated data, etc.
 - 307 • The authors should provide scripts to reproduce all experimental results for the new
308 proposed method and baselines. If only a subset of experiments are reproducible, they
309 should state which ones are omitted from the script and why.
 - 310 • At submission time, to preserve anonymity, the authors should release anonymized
311 versions (if applicable).
 - 312 • Providing as much information as possible in supplemental material (appended to the
313 paper) is recommended, but including URLs to data and code is permitted.

314 6. Experimental Setting/Details

315 Question: Does the paper specify all the training and test details (e.g., data splits, hyper-
316 parameters, how they were chosen, type of optimizer, etc.) necessary to understand the
317 results?
318

319 Answer: [Yes]

320 Guidelines:

- 321 • The answer NA means that the paper does not include experiments.
- 322 • The experimental setting should be presented in the core of the paper to a level of detail
323 that is necessary to appreciate the results and make sense of them.
- 324 • The full details can be provided either with the code, in appendix, or as supplemental
325 material.

326 7. Experiment Statistical Significance

327 Question: Does the paper report error bars suitably and correctly defined or other appropriate
328 information about the statistical significance of the experiments?

329 Answer: [Yes]

330 Guidelines:

- 331 • The answer NA means that the paper does not include experiments.
- 332 • The authors should answer "Yes" if the results are accompanied by error bars, confi-
333 dence intervals, or statistical significance tests, at least for the experiments that support
334 the main claims of the paper.
- 335 • The factors of variability that the error bars are capturing should be clearly stated (for
336 example, train/test split, initialization, random drawing of some parameter, or overall
337 run with given experimental conditions).
- 338 • The method for calculating the error bars should be explained (closed form formula,
339 call to a library function, bootstrap, etc.)
- 340 • The assumptions made should be given (e.g., Normally distributed errors).
- 341 • It should be clear whether the error bar is the standard deviation or the standard error
342 of the mean.
- 343 • It is OK to report 1-sigma error bars, but one should state it. The authors should
344 preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis
345 of Normality of errors is not verified.
- 346 • For asymmetric distributions, the authors should be careful not to show in tables or
347 figures symmetric error bars that would yield results that are out of range (e.g. negative
348 error rates).
- 349 • If error bars are reported in tables or plots, The authors should explain in the text how
350 they were calculated and reference the corresponding figures or tables in the text.

351 8. Experiments Compute Resources

352 Question: For each experiment, does the paper provide sufficient information on the com-
353 puter resources (type of compute workers, memory, time of execution) needed to reproduce
354 the experiments?

355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407

Answer: [Yes]

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [Yes]

Guidelines:

- 408 • The answer NA means that the paper poses no such risks.
- 409 • Released models that have a high risk for misuse or dual-use should be released with
- 410 necessary safeguards to allow for controlled use of the model, for example by requiring
- 411 that users adhere to usage guidelines or restrictions to access the model or implementing
- 412 safety filters.
- 413 • Datasets that have been scraped from the Internet could pose safety risks. The authors
- 414 should describe how they avoided releasing unsafe images.
- 415 • We recognize that providing effective safeguards is challenging, and many papers do
- 416 not require this, but we encourage authors to take this into account and make a best
- 417 faith effort.

418 12. Licenses for existing assets

419 Question: Are the creators or original owners of assets (e.g., code, data, models), used in
420 the paper, properly credited and are the license and terms of use explicitly mentioned and
421 properly respected?

422 Answer: [Yes]

423 Guidelines:

- 424 • The answer NA means that the paper does not use existing assets.
- 425 • The authors should cite the original paper that produced the code package or dataset.
- 426 • The authors should state which version of the asset is used and, if possible, include a
- 427 URL.
- 428 • The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- 429 • For scraped data from a particular source (e.g., website), the copyright and terms of
- 430 service of that source should be provided.
- 431 • If assets are released, the license, copyright information, and terms of use in the
- 432 package should be provided. For popular datasets, paperswithcode.com/datasets
- 433 has curated licenses for some datasets. Their licensing guide can help determine the
- 434 license of a dataset.
- 435 • For existing datasets that are re-packaged, both the original license and the license of
- 436 the derived asset (if it has changed) should be provided.
- 437 • If this information is not available online, the authors are encouraged to reach out to
- 438 the asset's creators.

439 13. New Assets

440 Question: Are new assets introduced in the paper well documented and is the documentation
441 provided alongside the assets?

442 Answer: [No]

443 Justification: In our approach, we build upon existing work by utilizing the BERT model,
444 incorporating the contextual information of both the head and tail entities. This allows us to
445 leverage richer contextual cues from both ends of the relationship, enhancing the model's
446 ability to understand and capture the nuances in the data.

447 Guidelines:

- 448 • The answer NA means that the paper does not release new assets.
- 449 • Researchers should communicate the details of the dataset/code/model as part of their
- 450 submissions via structured templates. This includes details about training, license,
- 451 limitations, etc.
- 452 • The paper should discuss whether and how consent was obtained from people whose
- 453 asset is used.
- 454 • At submission time, remember to anonymize your assets (if applicable). You can either
- 455 create an anonymized URL or include an anonymized zip file.

456 14. Crowdsourcing and Research with Human Subjects

457 Question: For crowdsourcing experiments and research with human subjects, does the paper
458 include the full text of instructions given to participants and screenshots, if applicable, as
459 well as details about compensation (if any)?

460 Answer: [NA]

461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.