
Prompt Learning Unlocked for App Promotion in the Wild

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 In recent times, mobile apps have increasingly incorporated app promotion ads to
2 promote other apps, raising cybersecurity and online commerce concerns related to
3 societal trust and recommendation systems. To effectively discover the intricate
4 nature of the app promotion graph data, we center around the graph completion
5 task, aiming to learn the connection patterns among diverse relations and entities.
6 However, accurately deciphering the connection patterns in such a large and
7 diverse graph presents significant challenges for deep learning models. To overcome
8 these challenges, we introduce `Prompt Promotion`, a transformer-based
9 framework that unlocks *prompt learning* capabilities by incorporating metapath-
10 and embedding-based prompts that provide valuable hints to guide the model’s
11 predictions for undetermined connection patterns. Experimental results show that
12 our `Prompt Promotion` model represents a pioneering prompt-based capability
13 in effectively completing the app promotion graph. It not only demonstrates
14 superior performance in heterogeneous graph completion in real-world scenarios,
15 but also exhibits strong generalization capabilities for diverse, complex, and noisy
16 connection patterns when paired with their respective prompts.

17 1 Introduction

18 Mobile applications, or apps, often incorporate
19 advertisements (ads) as a means of promotion
20 (Viennot et al., 2014; Liu et al., 2015), among
21 which app-promotion ads are commonly used by
22 Android app developers to promote other apps
23 (Research, 2023). However, concerns arise regarding
24 the trustfulness of the apps promoted through
25 these ads, given the competitive nature of the
26 industry and the potential for the promotion of
27 malicious apps (Rafieian and Yoganarasimhan, 2021;
28 Nath, 2012). Previous research has focused on
29 analyzing the behaviors of ad libraries within the
30 app promotion ecosystem (Grace et al., 2012;
31 Vallina-Rodriguez et al., 2012; Nath, 2015;
32 Jin et al., 2021; Liu et al., 2020). However, these
33 studies primarily examine the behaviors of ad
34 libraries themselves, and pay too little attention
35 to app propagation in terms of how massive
36 individuals exploit the app promotion ecosystem.
37 For instance, Figure 1 illustrates an app
promotion chain where a popular benign app
“Passport Photo Maker - ID/VISA” promotes a
greyware app “Photo Collage, Photo Editor”,
which in turn promotes malware “Flood-It!”,
a strategy game capable of scanning the local
network and stealing sensitive phone information.
Furthermore, these studies lack a comprehensive
understanding of app promotions, which involve
multiple heterogeneous actors beyond apps, such
as app markets, security vendors, and developers.



Figure 1: Example of malicious app promotion.

38 For example, Figure 2 provides an inference path to explain why an online messaging app, “Polish
 39 English Translation”, promotes “CallApp”. The underlying behaviors indicate that “Polish English
 40 Translation” shares the same developer as another translation app “Thai Chinese Translation”, which
 41 has been observed to promote “CallApp”. Hence, a more holistic approach that learns the intrinsic
 42 connection patterns among these various entities is necessary to deeply understand the complexities
 43 of the whole app promotion ecosystem and its implications for society and online commerce. To
 44 address these limitations, we employ insights of graph completion learning into the heterogeneous
 45 app promotion graph. Our goal is to predict unknown target entities based on known source entities
 46 and relation queries, thereby completing the full graph. By applying graph completion methods to
 47 the app promotion graph, we are able to learn representations that capture the intricate connection
 48 patterns among different types of entities and relations. This approach not only sheds light on the
 49 underlying dynamics of app promotion graphs, but also opens up possibilities for diverse applications.

50 Nevertheless, learning to complete the focused
 51 app promotion network is non-trivial, especially
 52 with datasets collected from the wild. Exist-
 53 ing methods for graph completion are either too
 54 simplistic for modeling the network complexity
 55 and information among relationships and enti-
 56 ties (Bordes et al., 2013; Sun et al., 2018; Yang




Figure 2: Example of inferred app promotion path.

57 et al., 2015), or they heavily rely on rich semantic information to train massive weight parameters
 58 (Wang et al., 2021; Lv et al., 2022; Yao et al., 2019), which contradicts the scarcity of semantic
 59 information in app promotion networks collected from the wild. Therefore, in this work, considering
 60 the challenge of modeling complex connection patterns while overcoming the limitations of existing
 61 techniques, we introduce our approach `Prompt Promotion`, which guides the model in learning
 62 the intricate connection patterns by incorporating a combination of embedding-based and metapath-
 63 based prompts. Leveraging the power of pretrained BERT, we design the embedding-based prompts,
 64 derived from pretrained embedding-based methods like DistMult (Yang et al., 2015), provide prior
 65 knowledge as hints to assist the model in making informed references. Additionally, we further craft
 66 metapath-based prompts by extracting not only valid but also informative metapaths for each queried
 67 relation. Subsequently, we combine the embedding-based and metapath-based prompts along with
 68 the query tokens, and randomly permute them to form the final input sequence for each query. The
 69 sequence is tokenized using the embedding-based method, replacing the original BERT tokenizer, to
 70 ensure that the tokens are projected into the same embedding space for the subsequent fine-tuning
 71 process. In summary, the contributions of this paper are:

- 72 • We propose a novel approach named `Prompt Promotion`, that addresses the challenge of
 73 modeling connection patterns in complex app promotion graphs by leveraging the pretrained BERT
 74 as the backbone model while incorporating both the embedding-based and metapath-based prompts
 75 to guide the model in learning the intricate patterns within the graph.
- 76 • We demonstrate the effectiveness of our approach through extensive experiments on our collected
 77 real-world dataset. The results show that our approach outperforms existing techniques in terms of
 78 accuracy and generalization capabilities in extracting diverse and complex connection patterns.
- 79 • We contribute to the research community by providing a deeper understanding of the app promo-
 80 tion ecosystem, its complexities, and implications for societal trust and online commerce. Our
 81 work sheds light on the potential applications of graph completion methods, specifically utilizing
 82 pretrained BERT, in improving trustworthiness in detecting malicious apps.

83 2 Background

84 In this section, we briefly introduce the definitions of a heterogeneous graph, metapath, and the task
 85 of heterogeneous graph completion, as well as the idea of prompt engineering and details about our
 86 app promotion graph dataset. We additionally refer related works in Appendix E.

87 2.1 Definitions

88 **Definition 1** (Heterogeneous Graph). A heterogeneous graph (HG) $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{X})$ consists of a node
 89 set \mathcal{V} , an edge set \mathcal{E} , and the optional features of the associated nodes and edges: $\mathcal{X} = (\mathcal{X}_{\mathcal{V}}, \mathcal{X}_{\mathcal{E}})$.

90 Each node’s type is mapped through the node type mapping function $\phi : \mathcal{V} \rightarrow \mathcal{A}$, and each edge’s
 91 type through the edge type mapping function $\psi : \mathcal{E} \rightarrow \mathcal{R}$, where \mathcal{A} and \mathcal{R} denotes the node and edge
 92 type set respectively. We represent each edge as a triple $(h, r, t) \in \mathcal{E}$ where $h, t \in \mathcal{V}, r \in \mathcal{R}$, and the
 93 edges are directional. For a heterogeneous graph, there exists the constrain $|\mathcal{A}| + |\mathcal{R}| > 2$.

94 **Definition 2** (Metapath). In a heterogeneous graph, a metapath represents a predefined sequence of
 95 node types and edge types that capture the desired semantic relationships between nodes. Formally, a
 96 metapath \mathcal{P} is denoted as $e_1 \xrightarrow{r_1} e_2 \xrightarrow{r_2} \dots e_L \xrightarrow{r_L} e_{L+1}$, where $r_i \in \mathcal{R}, e_i \in \mathcal{A}, r = r_1 \cdot r_2 \cdot \dots \cdot r_L$
 97 is the composite relation between entity type e_1 and e_{L+1} , and L is the length of the metapath.

98 **Definition 3** (Heterogeneous Graph Completion). For a valid query (h, r) where $h \in \mathcal{V}$ and $r \in \mathcal{R}$,
 99 a heterogeneous graph completion (HGC) task refers to discovering valid answers $\mathcal{T} \subset \mathcal{V}$ such that
 100 for all $t \in \mathcal{T}, (h, r, t) \in \mathcal{E}$.

101 2.2 Prompt Engineering

102 Prompt engineering is a systematic methodology widely employed in natural language processing
 103 applications to craft specific input signals to invoke desired output responses from machine learning
 104 models. Under the task of graph completion where input tokens are the queries, prompts can be
 105 designed as contextual semantic information related to the entities and relations, the query-related
 106 neighborhood, or other encoded information that guides the model to answer the query. Specifically
 107 for a query (h, r) in the graph completion task, the prompted input sequence is usually formulated as:

108
$$[<\text{bos}>] \text{ <prompts> } [<\text{sep}>] \text{ <h> <r> } [<\text{sep}>].$$

 109

110 2.3 App Promotion Dataset

111 2.3.1 Data Collection

112 The dataset pertaining to app promotion is gathered from three distinct perspectives. Initially, for
 113 each app, the package name, developer information, and category of each app are crawled from
 114 Google Play. Subsequently, an analysis is conducted on the app using VirusTotal to examine the
 115 flags associated with its security level, along with the corresponding URLs. Lastly, the manifest and
 116 signature of each app are inferred through the process of reverse engineering (e.g., interesting strings
 117 provided by the VirusTotal report). The promotion actions between apps are discovered by checking
 118 whether the clickable widgets in a UI from the source apps lead to the download page of the sink
 119 app. If so, then a `source_app <promotes> sink_app` relation is identified. The collected
 120 raw data is then used to construct the following HG for the capture of ample behavior patterns.

121 2.3.2 Graph Construction

122 In order to harness the informative attributes of applications, such as URLs and signatures, which
 123 are instrumental in forecasting elusive promotional strategies and discerning recurrent patterns in
 124 app promotion, we construct the App Promotion HG (APHG) to epitomize the sundry entities
 125 and relations inherent in the network. More details related to the entity statistics and relations of
 126 constructed graph are provided in Appendix A.

127 **Entities.** An APHG encapsulates distinct entities derived from the following app attributes: appli-
 128 cation package name, developer, application category, manifest, VirusTotal (VT) Engine, digital
 129 signature, and URL. The manifest entity encompasses app activities, providers, receivers, services,
 130 and permissions. Given the unique promotional behaviors demonstrated by benign, greyware, and
 131 malicious applications, we further classify the application package name into these three discrete
 132 classes, and extend the aggregate count of entity types within our framework to nine.

133 **Relations.** We consider multiple directional relations among the entities defined above to capture
 134 their interactive behaviors: *app-promote-app*, *app-include-signature*, *engine-detect-app*, *app-belong-*
 135 *category*, *developer-involve-category*, *developer-develop-app*, *app-access-URL*, *developer-use-URL*,
 136 *app-own-manifest*. Since apps with different security levels follow different behavior patterns, we
 137 further divide them into three sub-classes: *benign*, *grey* and *malicious*. In total, the above relations
 138 are extended to twenty-nine classes of relation types. Note that all the relations are directional, and
 139 each query only associates with one of the constructed directional relations, excluding the reverse
 140 relations. Despite the potential to gather additional information, neither the entities nor the relations
 141 are associated with any features. Therefore, our APHG is denoted as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$.

142 **2.3.3 Task Motivation**

143 Despite the relations in place abundantly capturing intrinsic application information, instances of
 144 information paucity are far from scarce. This scarcity impedes our ability in gathering all relevant
 145 information (see an illustration on Google Play in Appendix B), which in turn substantially impedes
 146 our capacity to decipher patterns in application promotion behavior. Thus, to alleviate the burden of
 147 information scarcity, we propose to first target the HGC task on our app promotion graph.

148 **3 Unlocking Prompt Definition on HGC**

149 **3.1 Overall Framework**

150 Our Prompt Promotion approach leverages the pre-trained BERT (Devlin et al., 2019) as the
 151 transformer encoder to encode the tokenized input sequence of each query. We use an aggregator
 152 to consolidate the output sequence and a two-layer MLP as the prediction head to perform the final
 153 task. The overall framework is depicted in Figure 3. We incorporate this design for two reasons:
 154 (1) encoding the query instead of the triple mitigates the calculation overhead when responding to a
 155 specific query, and (2) the attention mechanism in BERT assigns global attention to the provided input,
 156 including the designed prompts. We extend the input for each query into three parts: embedding-based
 157 prompts, metapath-based prompts, and the query itself, consisting of a source-relation pair. In the
 158 following content, we provide details regarding the two sets of prompts.

159 **3.2 Embedding-based Prompts**

160 Prior embedding-based models have demon-
 161 strated remarkable performance on various pub-
 162 lic benchmark datasets, making them state-of-
 163 the-art solutions. These models possess inher-
 164 ent simplicity that renders them proficient tok-
 165 enizers, effectively mapping entity and relation
 166 tokens to a shared semantic space. In this pa-
 167 per, we select DistMult (Yang et al., 2015) as
 168 the pre-trained embedding-based method to tok-
 169 enize the entity and relation tokens. Note that
 170 this is a designer’s choice and can be substituted
 171 with any other methods that fit our framework.
 172 The n embedding-based prompts are defined as
 173 the top- n predicted entities by the pretrained
 174 embedding-based methods according to the pre-
 175 dicted scores, denoted as C^e . These prompts
 176 serve as prior knowledge that assists the model
 177 in making informed references. For instance,
 178 when considering the query “which app does
 179 Instagram promote?”, we provide additional
 180 prompts in the form of a hint, such as “I am not 100% sure, but I believe these apps might be
 181 the answers.” This supplementary information aids the model in further generating more accurate
 182 and contextually relevant responses. The filtered prompted entities are then tokenized with the
 183 corresponding embeddings learned by the pre-trained embedding-based method.

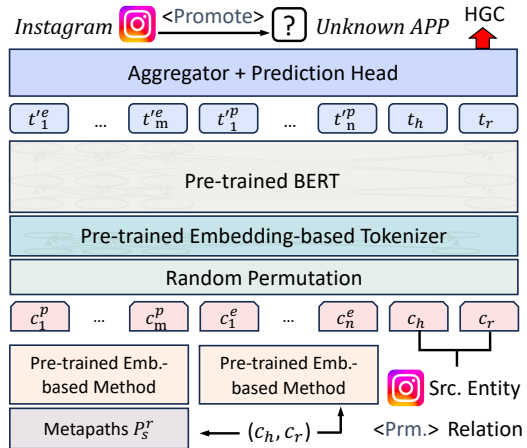


Figure 3: Overall framework of our Prompt Promotion.

184 **3.3 Metapath-based Prompts**

185 While embedding-based prompts serve as hints from the pioneers, they may neglect true answers due
 186 to their narrowed perspectives. Typically, embedding-based methods utilize geometric operations
 187 in the representation space, resulting in prompts that share similarities from a geometric perspec-
 188 tive. Although entities in a knowledge graph inherently possess semantic meanings, we posit that
 189 the semantic information of entities in a heterogeneous graph can be alternatively extracted from
 190 metapaths. Metapaths offer a means to capture and encode meaningful relationships within the
 191 graph, facilitating the extraction of valuable insights. Therefore, we further provide the model with
 192 prompts from another perspective, i.e., the metapath-based prompts from the semantic perspective.

193 The key assumption lies in the connection between a certain metapath and the queried relation. In the
 194 following content, we first introduce the measure of the correlation between a metapath and a query,
 195 and then illustrate how to utilize the correlation to create the metapath-based prompts.

196 3.3.1 Query-Metapath Correlation

197 For a clearer clarification, we first define the functions $src(\cdot)$ and $dst(\cdot)$ as the source and destination
 198 entity type extractions for a relation r respectively. Regarding a specific queried relation, we make
 199 the following definition:

200 **Definition 4** (*r*-valid Metapath). *A metapath $p = e_1 \xrightarrow{r_1} e_2 \xrightarrow{r_2} \dots e_L \xrightarrow{r_L} e_{L+1}$ is *r*-valid if and*
 201 *only if e_1 and e_{L+1} are the source and destination entity types of relation r , respectively.*

202 For example, for the relation $r = benign-access-URL$, a corresponding valid metapath includes but is
 203 not limited to $benign \xleftarrow{develop} developer \xrightarrow{use} URL$, where $src(r) = benign$ and $dst(r) = URL$. The
 204 first step of linking a queried relation with a certain metapath is to identify all the *r*-valid metapaths.
 205 For multi-hop reasoning tasks, the answers to a query usually lie within three hops. We control the
 206 length of the metapath as $L \leq 2$ and conduct an exhaustive search for each query relation $r \in \mathcal{R}$,
 207 where \mathcal{R} denotes the set of all relations. The set of all *r*-valid metapaths is denoted as \mathcal{P}_r . Note that
 208 when searching for *r*-valid metapaths, we also consider the reverse relation of the original relation,
 209 since there exist entities with only outgoing edges, and reverse relations do not change the semantic
 210 meanings. However, we only consider the original relations as the queried relations. The metapaths
 211 in \mathcal{P}_r are valid, but not necessarily informative. In other words, \mathcal{P}_r does not inform us how relevant
 212 each $p \in \mathcal{P}_r$ is to r . To quantify the correlation, we make the following definitions:

213 **Definition 5** (*p*-Hit). *For a specific triple (h, r, t) , where r is the relation, h is the source entity*
 214 *such that $h \in \mathcal{H} \subset \mathcal{V}$ and $\phi(h) = src(r)$, t is the destination entity such that $t \in \mathcal{T} \subset \mathcal{V}$ and*
 215 *$\phi(t) = dst(r)$, we say the triple (h, r, t) is *p*-Hit if and only if there exist at least one path from h to t*
 216 *such that this path is an instance of the metapath p .*

217 **Definition 6** (*p*-Hit Ratio). *For a specific triple (h, r, t) , if this triple is *p*-Hit, then the *p*-hit ratio α*
 218 *of this triple is defined as the ratio of t among all other entities reached by the metapath p ; otherwise,*
 219 *the *p*-hit ratio of this triple is zero.*

220 **Definition 7** (*r-p* Ratio). *For a specific relation r , a metapath $p \in \mathcal{P}_r$, and all true (h, r, t) triples,*
 221 *the corresponding *r-p* ratio is defined as the averaged hit ratio of all true r related triples, i.e.,*
 222 *triples constructed with relation r . Note that the ratio is calculated based on a filtered setting: if*
 223 *t' is a correct answer to the query (h, r) when evaluating on the answer t , we remove t from the*
 224 *denominator.*

225 We here provide a concrete example for exemplification. Consider the relation *benign-access-URL*
 226 and its valid metapath $p = benign \xleftarrow{develop} developer \xrightarrow{use} URL$. For each true triple $(h, benign-$
 227 $access-URL, t)$ such that $\phi(h) = benign$ and $\phi(t) = URL$, denote the set of accessible URLs to the
 228 query $(h, benign-access-URL)$ as \mathcal{T}_h , and the set of all URL entities reached by following metapath
 229 p starting from h as \mathcal{T}_h^p . If the triple is *p*-Hit, then the hit ratio is calculated as $\alpha = 1/(|\mathcal{T}_h^p \setminus \mathcal{T}_h| - 1)$;
 230 otherwise, $\alpha = 0$. We minus one in the denominator because $t \in \mathcal{T}_h$. The *r-p* ratio of the relation
 231 *benign-access-URL* is then calculated as the averaged α of all the related true triples. Naturally, if
 232 a metapath p is highly correlated with r for a specific source entity h , the corresponding α should be
 233 high. We utilize the *r-p* ratio of each relation-metapath pair as the correlation indicator to select the
 234 top- m metapaths for further prompt generation, and denote the m selected metapaths as \mathcal{P}_r^s .

235 3.3.2 Metapath-based Prompt Generation

236 Even though we select m metapaths for each query, some metapaths may contain noise. This is
 237 especially true when a metapath reaches a high-degree entity, resulting in a significant expansion
 238 of the candidate pool. In such cases, these prompts may not provide any substantial additional
 239 information beyond what is already known, rendering them less informative. To address this, we
 240 apply a candidate filtering method. Specifically, we utilize a limit l to separate metapaths that lead to
 241 large or small candidate sizes. For small-sized candidates, we perform the *union* operation, and for
 242 large-sized candidates, we perform the *intersect* operation. The rationale is as follows: some queries
 243 may not be highly relevant to just one metapath, in which case the number of candidates is usually
 244 large, and we rely on the *intersect* operation to filter out noise. On the other hand, some queries may

245 be explained by more than one metapath, in which case the size of the candidate pool is usually small,
 246 and the *union* operation considers all conditions.

247 After the filtering process, we empiri-
 248 cally evaluate the correlation between one
 249 queried relation and the filtered candidates.
 250 Particularly, We calculate the average size
 251 of the filtered candidate pools s_r for all r
 252 related triples, as well as the hit ratio h_r of
 253 the correct answer for each type of query
 254 among the candidate pools. In addition, we
 255 denote the base hit ratio as $b_r = s_r/|\phi(t)|$
 256 and the magnification as $m_r = h_r/b_r$.
 257
 258
 259
 260
 261
 262
 263
 264
 265

Table 1: Empirical evaluation results of the correlation between metapath and relation.

Relation	h_r	s_r	b_r	m_r
mal-belong-category	0.922	4.932	0.137	6.732
benign-access-URL	0.879	129.329	0.007	128.1
developer-use-URL	0.457	42.905	0.002	200.486
grey-promote-grey	0.660	154.786	0.135	4.876

Table 1 presents a selection of the evaluation results due to the large size of \mathcal{R} . The table provides rich information: (1) the selected metapaths for some relations are highly correlated with their relations, indicated by high h_r and low s_r (e.g., mal-belong-category); (2) some other relations provide a considerable amount of correlation, indicated by a large m_r , but may lead to a high hit ratio (e.g., benign-access-URL) or a low hit ratio (e.g., developer-use-URL), affected by s_r ; (3) there are also cases in the middle with decent h_r and s_r (e.g., grey-promote-grey). Nevertheless, the results confirm that metapaths provide information regarding the query, regardless of high or low h_r . To reduce the size of the input prompts, we further utilize an embedding-based method to select the top- m prompts among the candidate set C_h^r as the final metapath-based prompts, denoted as C^p .

266 3.4 Combined Input Sequence

267 For a query (h, r) , we concatenate the embedding-based prompts C^e , the metapath-based prompts
 268 C^p , and the query token h and r as the final input sequence. Before feeding the constructed sequence
 269 into the pre-trained BERT model, we randomly permute the tokens. This step is essential in forcing
 270 the BERT model to learn the intrinsic connection between the query and the answer, rather than
 271 relying too much on the prompts. We validate the necessity of this step in the following experiments.
 272 After the permutation, the input sequence is tokenized via the embedding-based method, replacing
 273 the original BERT tokenizer. Finally, we adopt the binary cross entropy loss for the HGC task. We
 274 provide the pseudo code for our method in Appendix C.

275 4 Experiment

276 4.1 Setup

277 We test our method’s effectiveness over the constructed APHG as described in Section 2.3. For
 278 comparison, we carefully select DistMult (Yang et al., 2015), ComplEX (Trouillon et al., 2016),
 279 ConvE (Dettmers et al., 2018), HittER (Chen et al., 2021), and LTE (Zhang et al., 2022) as the
 280 baselines, for they can be easily adapted to our HGC task. For evaluation purposes, we adopt two
 281 key metrics: mean reciprocal rank (MRR) and Hits@K, and higher values of MRR and Hits@K
 282 indicate better performance in accurately ranking and identifying the correct candidates in the graph
 283 completion task. We use a pre-trained DistMult (Yang et al., 2015) as the backbone model to tokenize
 284 the entities and relations as low-dimensional vectors, and utilize a pre-trained ComplEX (Trouillon
 285 et al., 2016) for prompt filtering. Note that these choices are a matter of preference, and can be
 286 substituted with other embedding-based methods such as TransE (Bordes et al., 2013). We consider
 287 two settings under our framework: *w/ Rand. Perm.* denotes that we randomly permute the input
 288 tokens before the encoding process, and *w/o Rand. Perm.* suggests otherwise. The input sequence
 289 is decomposed into three essential components - the embedding-based prompts, metapath-based
 290 prompts, and the query. Based on the above settings and components, we define model variants as
 291 shown in Table 2. More detailed experimental setups are provided in Appendix D due to space limit.

292 4.2 Performance on App Promption

293 The performance comparison in Table 3 demonstrates that our model outperforms the other baselines
 294 by a significant margin. This improvement can be attributed to two key factors: the incorporation
 295 of the designed prompts and the utilization of random permutation. While our model utilizes




Figure 4: Results of component-differed variants, including ours (Full-Prompt w/ Rand. Perm.).

296 DistMult (Yang et al., 2015) as the backbone, it extends its capabilities beyond a simple multiplication
 297 projection of the queried source entity and relation embeddings. This is evident from the consistent
 298 notable performance enhancement achieved by our model. We also observe that as the value of
 299 K increases, the performance gap between our model and the baselines gradually diminishes. We
 300 hypothesize that our model follows a two-step inference process: first, it processes the provided
 301 prompts and attempts to identify potential answers out of the input sequence. If the correct answers
 302 are present in the prompts, the model can recognize them with relatively high probabilities, leading to
 303 higher hit ratios when K is small. This aspect of the task is relatively straightforward. However, if
 304 the answers are not found in the provided prompts, the model transits to another task and endeavors
 305 to generate an answer by considering all the given hints. This second task tests the model’s ability to
 306 deduce query patterns and is inherently more challenging. We refer to this hypothesis as the “dual-
 307 task” hypothesis, which suggests that our model performs and excels at both the answer identification
 308 and answer generation tasks. Additionally, we observe a notable performance downgrade among all
 309 the variants compared to the best. Under most conditions, the BERT encoder significantly improves
 310 Hit@1 performance, suggesting that our framework focuses more on direct query answering, rather
 311 than pattern matching. We provide more detailed analysis in the following section to validate our
 312 “dual-task” hypothesis, and examine the model’s capabilities under several conditions.

313 4.3 Component Analysis for Prompt Designs

314 In this part, we further analyze the
 315 impacts of each component in our
 316 framework to confirm the necessity
 317 of constructing our model as de-
 318 signed, as well as providing support-
 319 ive evidence for our “dual-task” hy-
 320 pothesis. We add another variant
 321 *Random-Prompt*, where the input se-
 322 quence is constructed with randomly sampled prompts plus the query tokens.

Table 2: Definitions of variants of our Prompt Promotion.

Variant	Emb. Prm.	Mtp. Prm.	Query	Rand. Perm.
Base	✗	✗	✓	✓
Emb.-based Only	✓	✗	✓	✓
Mtp.-based Only	✗	✓	✓	✓
Ours w/o Rand. Perm.	✓	✓	✓	✗
Ours (Prompt Promotion)	✓	✓	✓	✓

323 4.3.1 Performance Comparison

324 The performance of the variants is shown in Figure 4. Note that we skip the *w/ Rand. Perm.* setting
 325 for the *Base* variant since the order of two tokens is trivial and randomly permuting them does not
 326 affect the performance too much. From Figure 4, we make the following key observations:

- 327 • We consider the *Base* variant as training the BERT encoder to replace the matrix multiplication
 328 operation in DistMult. While it does not induce model collapse, it is still challenging to enforce a
 329 BERT encoder to fill the role of the operation. This observation inspires our Prompt Promotion
 330 approach, which detours the functionality replication of matrix multiplication and extends the
 331 power beyond it by introducing additional prompts.
- 332 • The addition of randomly generated prompts completely collapses the model, regardless of the use
 333 of random permutation. This is because the model is overwhelmed with not only the HGC task, but
 334 also the identification of the queried entity and relation tokens. This suggests the requirements of
 335 carefully crafted prompts with very limited noises.

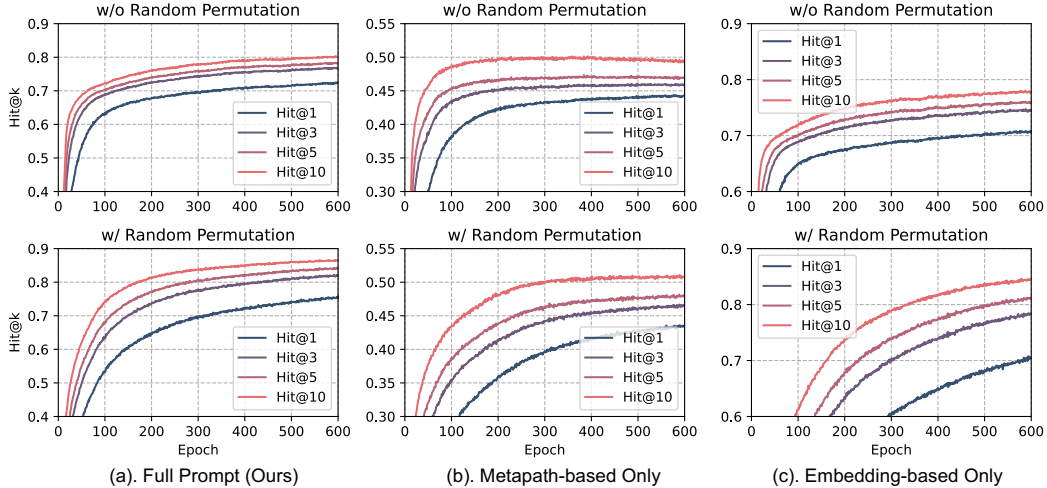


Figure 5: Learning dynamics of models with full prompts, metapath-based prompts only, and embedding-based prompts only.

- 336 • The *Embedding-based Only* variant yields decent performance under the two settings, especially
 337 for the hit ratios with small K 's. This not only validates the necessity of the embedding-based
 338 prompts, but also confirms one side of the hypothesis - the BERT structure is considerably good at
 339 identifying the existing answer among the input prompts.
- 340 • The fact that *Metapath-based Only* underperforms the *Base* can also be explained by the unavoidable
 341 noise introduced in the prompts. In comparison, although *Embedding-based Only* also takes extra
 342 prompts, these prompts are structurally similar in the embedding space, while the noise introduced
 343 by merely following the metapaths is intractable.
- 344 • *Full-Prompt* outperforms all other variants under the two settings, suggesting the necessity in the
 345 combination of the two sets of prompts. We also discover that as K increases, the gap between our
 346 variants and the baselines decreases faster under the *w/o Rand. Perm.* setting, compared with the
 347 other. This is because the model relies too much on identifying the existing prompts by splitting
 348 less explanation power in deducing the query patterns. Randomly permuting the input tokens
 349 mingles the prompts all together, therefore forcing the model to focus on the intrinsic connection
 350 between the prompts and the query, rather than the one hooked by the token positions.

351 4.3.2 Training Dynamics Analysis

352 We analyze the train-
 353 ing dynamics of the
 354 *Embedding-based Only*,
 355 *Metapath-based Only*, and
 356 *Full-Prompt* variants under
 357 two settings with their
 358 learning curves shown in
 359 Figure 5. Comparing from
 360 the setting perspective,
 361 we observe that models
 362 converge slower under
 363 *w/ Rand.Perm.*. This is
 364 because variants under *w/o*
 365 *Rand.Perm.* tends to take
 366 the shortcut solution by
 367 memorizing the positions,
 368 rather than learning the behavior patterns. Identifying the shortcut token's positions, compared
 369 with the HGC task, is a relatively easier task that requires less model complexity and learning time.
 370 This aligns with our "dual-task" hypothesis - the easier line of task is to identify the answer from
 371 the prompts, leading to faster convergence, and the harder one is to deduce the query patterns,
 372 corresponding to a relatively slower convergence. Additionally, we find that the gaps in hit ratios for

Table 3: Performance comparison with the baselines. Best results are bolded, and runner-ups are underlined.

Model	Hit@1	Hit@3	Hit@5	Hit@10	MRR
DisMult (Yang et al., 2015)	.6040	.7280	.7550	.8350	.6840
ComplEX (Trouillon et al., 2016)	.6680	.7780	.8180	.8650	.7370
ConvE (Dettmers et al., 2018)	.6400	.7460	.7950	.8490	.7110
HittER (Chen et al., 2021)	.5505	.6758	.7227	.7862	.6312
ConvE-LTE (Zhang et al., 2022)	.6350	.7444	.7918	.8506	.6602
Distmult-LTE (Zhang et al., 2022)	.6381	.7651	.8083	<u>.8677</u>	.7174
Base	.7246	.7610	.7729	.7895	.7481
Emb.-based Only	<u>.7786</u>	<u>.8272</u>	<u>.8447</u>	.8672	<u>.8096</u>
Mtp.-based Only	.4567	.4740	.4843	.5082	.4795
Ours w/o. Rand. Perm.	.7383	.7817	.7940	.8118	.7653
Ours	.8393	.8710	.8802	.8922	.8587

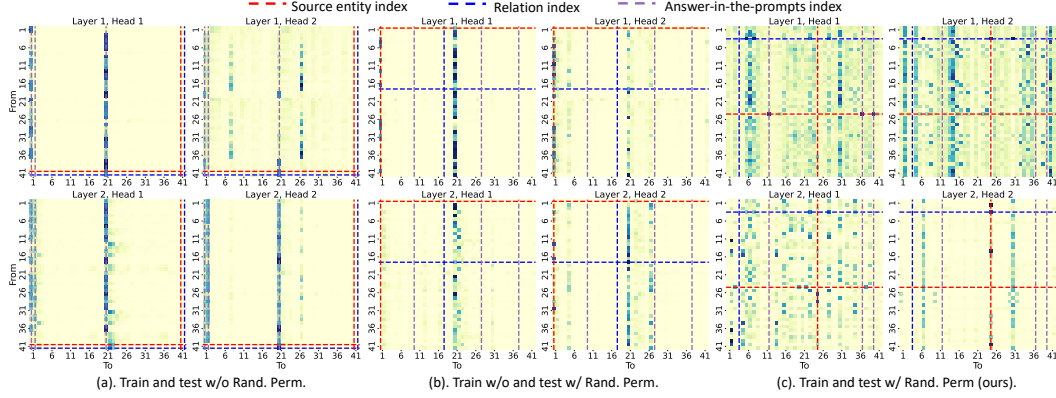


Figure 6: Attention heatmaps of the case under three settings.

373 different K 's are larger under *w/ Rand. Perm.*, indicating better generality and pattern extrapolation
 374 abilities. Among the variants, *Full-Prompt* exhibits reasonable learning behavior. It avoids saturating
 375 too quickly like *Metapath-based Only* due to less introduced noise, and does not take excessively
 376 long to achieve performance improvement like *Embedding-based Only*, which relies heavily on
 377 accessible shortcuts that hinder generality.

378 4.4 Random Permutation on Model Learning

379 To further analyze how the random permutation affects the model learning, we empirically study the
 380 model behavior under a specific query case. Consider the *Full-Prompt* variant, where we differ the
 381 train and test conditions: (a) We train and test the variant under *w/o. Rand. Perm.*; (b) We train the
 382 variant *w/o. Rand. Perm.*, but test it under *w/ Rand. Perm.*; (c) We train and test the variant under *w/
 383 Rand. Perm.* Regarding a specific query, we show the normalized attention scores heat map under the
 384 three conditions in Figure 6. The rankings of the correct answer under the three conditions are 1, 133,
 385 and 1 respectively. Under condition (a), we see the model consistently pay heavy attention to tokens
 386 on positions 1 and 21. This is because we set $m=20$ and $n=20$, and the most probable answers can
 387 usually be found in these positions. Without random permutation, the model quickly identifies the
 388 shortcut, rather than paying extra attention to the query (indexed by the red and blue dotted lines).
 389 Under condition (b), the model failed to assign a high ranking to the correct answer. Due to the
 390 random permutation, tokens on positions 1 and 21 no longer provide precise information as the
 391 model assumes, making it overwhelmed with the introduced randomness. This can also be confirmed
 392 with small attention scores assigned to the query and the potential answers. Therefore, randomly
 393 permuting the input sequence acts as a potential and effective attack to variants trained under *w/o
 394 Rand. Perm.* The model trained and tested under *w/ Rand. Perm.* as we designed, on the other
 395 hand, assigns much more even attention to the input sequence. More specifically, it learns to assign
 396 attention to the potential answers in the input (red and purple line intersections in Layer 1, Head 1),
 397 the source entity (red vertical dotted line in Layer 2, Head 2), as well as other important information
 398 in deems important (tokens indexed by 7, 31, etc.). This confirms that random permutation enhances
 399 the model's ability to learn the intrinsic connection between the query and the answer, reducing
 400 reliance on input prompts and increasing robustness and generality.

401 5 Conclusion

402 In this work, we focus on the heterogeneous graph completion task in the context of app promotion,
 403 and propose a prompt-based approach named `Prompt Promotion` that leverages a pre-trained
 404 BERT to model the connection patterns in the complex app promotion ecosystem. Specifically, by
 405 incorporating both embedding-based and metapath-based prompts, our model first unlocks the prompt
 406 learning for app promotion graphs, and achieves superior performance compared to baselines. In
 407 addition, we conduct thorough analysis regarding the components, training dynamics to illustrate
 408 the delicacy of our designed framework. The contributions of this research include advancing the
 409 understanding of app promotion networks, improving trustworthiness in recommender systems, and
 410 detecting promotion traces of malicious apps. Future directions involve exploring additional prompt
 411 generation strategies and further enhancing the model's performance.

412 References

- 413 Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013.
414 Translating embeddings for modeling multi-relational data. In *Advances in neural information*
415 *processing systems*. 2, 6, 14
- 416 Sanxing Chen, Xiaodong Liu, Jianfeng Gao, Jian Jiao, Ruofei Zhang, and Yangfeng Ji. 2021.
417 HittER: Hierarchical Transformers for Knowledge Graph Embeddings. In *Conference on Empirical*
418 *Methods in Natural Language Processing*. 6, 8, 13, 14
- 419 Tim Dettmers, Minervini Pasquale, Stenetorp Pontus, and Sebastian Riedel. 2018. Convolutional 2D
420 Knowledge Graph Embeddings. In *AAAI Conference on Artificial Intelligence*. 6, 8, 13, 14
- 421 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training
422 of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019*
423 *Conference of the North American Chapter of the Association for Computational Linguistics:*
424 *Human Language Technologies*. Association for Computational Linguistics. 4
- 425 Michael C Grace, Wu Zhou, Xuxian Jiang, and Ahmad-Reza Sadeghi. 2012. Unsafe exposure
426 analysis of mobile in-app advertisements. In *ACM conference on Security and Privacy in Wireless*
427 *and Mobile Networks*. 1
- 428 Michaela Hardt and Suman Nath. 2012. Privacy-aware personalization for mobile advertising. In
429 *ACM conference on Computer and communications security*. 1
- 430 Ling Jin, Boyuan He, Guangyao Weng, Haitao Xu, Yan Chen, and Guanyu Guo. 2021. MAdLens:
431 Investigating into Android In-App Ad Practice at API Granularity. *IEEE Transactions on Mobile*
432 *Computing* (2021). 1
- 433 Bin Liu, Bin Liu, Hongxia Jin, and Ramesh Govindan. 2015. Efficient privilege de-escalation for ad
434 libraries in mobile apps. In *Annual International Conference on Mobile systems, Applications, and*
435 *Services*. 1
- 436 Tianming Liu, Haoyu Wang, Li Li, Xiapu Luo, Feng Dong, Yao Guo, Liu Wang, Tegawendé
437 Bissyandé, and Jacques Klein. 2020. MadDroid: Characterizing and detecting devious ad contents
438 for android apps. In *The Web Conference*. 1
- 439 Xin Lv, Yankai Lin, Yixin Cao, Lei Hou, Juanzi Li, Zhiyuan Liu, Peng Li, and Jie Zhou. 2022. Do
440 pre-trained models benefit knowledge graph completion? a reliable evaluation and a reasonable
441 approach. In *Findings of the Association for Computational Linguistics*. 2, 14
- 442 Suman Nath. 2015. Madscope: Characterizing mobile in-app targeted ads. In *Annual International*
443 *Conference on Mobile Systems, Applications, and Services*. 1
- 444 Omid Rafeian and Hema Yoganarasimhan. 2021. Targeting and privacy in mobile advertising.
445 *Marketing Science* (2021). 1
- 446 Google Research. 2023. How people discover, use, and stay engaged with apps. *Think with Google*
447 (2023). 1
- 448 Soel Son, Daehyeok Kim, and Vitaly Shmatikov. 2017. What Mobile Ads Know About Mobile
449 Users. Internet Society. 1
- 450 Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2018. RotatE: Knowledge Graph
451 Embedding by Relational Rotation in Complex Space. In *International Conference on Learning*
452 *Representations*. 2, 14
- 453 Théo Trouillon, Johannes Welbl, Sebastian Riedel, Eric Gaussier, and Guillaume Bouchard. 2016.
454 Complex Embeddings for Simple Link Prediction. In *International Conference on Machine*
455 *Learning*. 6, 8, 13, 14
- 456 Narseo Vallina-Rodriguez, Jay Shah, Alessandro Finamore, Yan Grunenberger, Konstantina Papa-
457 giannaki, Hamed Haddadi, and Jon Crowcroft. 2012. Breaking for commercials: characterizing
458 mobile advertising. In *Internet Measurement Conference*. 1

- 459 Nicolas Viennot, Edward Garcia, and Jason Nieh. 2014. A measurement study of google play. In
460 *ACM international conference on Measurement and modeling of computer systems*. 1
- 461 Bo Wang, Tao Shen, Guodong Long, Tianyi Zhou, Ying Wang, and Yi Chang. 2021. Structure-
462 augmented text representation learning for efficient knowledge graph completion. In *The Web*
463 *Conference*. 2, 14
- 464 Xin Xie, Ningyu Zhang, Zhoubo Li, Shumin Deng, Hui Chen, Feiyu Xiong, Mosha Chen, and Huajun
465 Chen. 2022. From discrimination to generation: knowledge graph completion with generative
466 transformer. In *The Web Conference*. 14
- 467 Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2015. Embedding Entities and
468 Relations for Learning and Inference in Knowledge Bases. In *3rd International Conference on*
469 *Learning Representations, ICLR*. 2, 4, 6, 7, 8, 13, 14
- 470 Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. KG-BERT: BERT for knowledge graph comple-
471 tion. *arXiv preprint arXiv:1909.03193* (2019). 2, 14
- 472 Zhanqiu Zhang, Jie Wang, Jieping Ye, and Feng Wu. 2022. Rethinking Graph Convolutional Networks
473 in Knowledge Graph Completion. In *The Web Conference*. 6, 8, 13

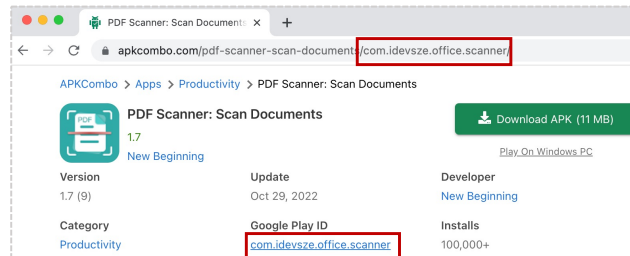
474 A App Promotion Heterogeneous Graph

475 Table 4 shows the statistics of the entities in the constructed App Promotion Heterogeneous Graph (APHG).
 476
 477 In addition, we define the relations as follows: (1) R1: an
 478 app-promote-app relation indicates that there exists a promotion link from the subject app
 479 to the object app; (2) R2: an
 480 app-include-signature relation means that a digital signature can be used to verify the
 481 authenticity and integrity of the app package; (3) R3: an engine-detect-app relation indi-
 482 cates that a VT engine marks an app with a specific flag (e.g., adware or Trojan); (4) R4: an
 483 app-belong-category relation represents that an app belongs to a specific app category categorized by
 484 Google Play; (5) R5: a developer-involve-category relation suggests that an app created
 485 by the developer is categorized into a specific app category; (6) R6: a developer-develop-app
 486 relation signifies that a developer develops an app; (7) R7: an app-access-URL relation denotes
 487 that an app has access to a specific URL; (8) R8: a developer-use-URL relation indicates that the
 488 app developed by the developer may access a specific URL; (9) R9: an app-own-manifest rela-
 489 tion represents that an app is associated with a specific manifest file. Since apps with different security
 490 levels follow different behavior patterns, we further divide the apps into three classes. For example, the
 491 relation app-belong-category is extended to three relations: benign-belong-category,
 492 grey-belong-category, and mal-belong-category. As a result, the above relations are
 493 extended to twenty-nine classes of relation types.
 494
 495
 496
 497


Table 4: Numbers and types for entities (or nodes).

type	Signature	VT Engine	Category	Developer	URL
num.	185	65	36	3139	18870
type	Manifest	Benign	Greyware	Malware	Total
num.	10269	3961	1143	363	38031

498 B Illustration of Information Scarcity



(a). APKCombo shows that the app was available on Google Play



(b). Google Play server cannot find the app anymore

Figure 7: Illustration of information scarcity on Google Play.

499 We here provide an illustration of information scarcity of within the app promotion ecosystem. As
 500 illustrated in Figure 7, the app *PDF Scanner*, which acts as a seed app and plays an instrumental role
 501 in promoting subsequent apps, was once available on Google Play. However, by relying solely on
 502 Google Play as our information source, we inevitably encounter instances where certain attributes,
 503 such as those related to the developer, are absent. Such omissions of information substantially impede
 504 our capacity to decipher patterns in application promotion behavior, and further motivate us to target
 505 the HGC task on our app promotion graph.

Algorithm 1 Prompt Promotion: a simplified PyTorch-style Pseudocode of our method on the HGC task.

```
# model: BERT-based model
# pretrained_kge: pretrained KGE method
# M: filtered metapaths for each relation
# Train model for N epochs
for query, target in dataloader:
    # Obtain emb-based prompts
    emb_prompt = pretrained_kge(query)[:n]

    # Find all reachable entities
    reached_ent = follow_metapath(query, M)

    # Sample k metapath-based prompts
    mtp_prompt = sample(reached_ent, m)

    # Forward
    input_seq = rand_perm(concat(emb_prompt, mtp_prompt, query))
    pred = model(input_seq)
    loss = CrossEntropyLoss(pred, target)

    # Optimize model with loss backward
    loss.backward()
    optimizer.step()
```

506 C Pseudo Code for Prompt Promotion

507 We provide the PyTorch style pseudocode of our proposed Prompt Promotion in Alg. 1 over the
508 app promotion HGC task.

509 D Experimental Setups

510 D.0.1 Dataset

511 Our app promotion dataset is collected from AndroZoo, a well-maintained and regularly updated
512 repository that provides various versions of apps from official app markets like Google Play. The
513 dataset encompasses apps released between January 1st, 2018, and February 3rd, 2023. We classify
514 the apps into three categories based on the number of engines that flag them on VirusTotal. Malware
515 apps are flagged by at least 10 engines, greyware apps are flagged by 1 to 9 engines, and benign apps
516 are not flagged by any engine on VirusTotal. Our seed dataset comprises approximately 48,000 apps,
517 evenly distributed among the three classes, providing a diverse set of apps representing different
518 levels of potential security risks. More details regarding the dataset and the construction for APHG
519 are provided in Section 2.3.

520 D.0.2 Baselines

521 We compare our approach against several baseline models commonly used in the graph completion
522 task:

- 523 • **DistMult** Yang et al. (2015): DistMult represents entities and relations as low-dimensional vectors
524 and utilizes a bilinear dot product scoring function for link prediction.
- 525 • **Complex** Trouillon et al. (2016): Complex extends DistMult by using complex-valued em-
526 beddings, allowing for a more expressive representation and remaining linear in both space and
527 time.
- 528 • **ConvE** Dettmers et al. (2018): ConvE employs a convolutional neural network architecture to
529 encode entities and relations. It operates on 2D tensors to capture local patterns and dependencies
530 within the knowledge graph.
- 531 • **HittER** Chen et al. (2021): HittER utilizes hierarchical transformers to learn knowledge graph
532 embeddings, balancing the contextual relational information and the information from the training
533 entity.
- 534 • **LTE** Zhang et al. (2022): LTE extends embedding-based methods by equipping existing knowl-
535 edge graph embedding models with linearly transformed entity embeddings. It mines semantic

536 information from entity representations to enhance the model performance. In this paper, we select
537 DistMult and ConvE as the backbones, denoted as DistMult-LTE and ConvE-LTE respectively.

538 **D.0.3 Evaluation Metrics**

539 We evaluate the graph completion performance using two key metrics: mean reciprocal rank (MRR)
540 and Hits@K. We empirically set the beam size for MRR as 256.

541 **D.0.4 Implementation Details**

542 We use a pre-trained DistMult [Yang et al. \(2015\)](#) as the backbone model to tokenize the entities and
543 relations as low-dimensional vectors. Note that this choice is a matter of preference, and can be substi-
544 tuted with other embedding-based methods such as TransE [Bordes et al. \(2013\)](#). CompLEX [Trouillon
545 et al. \(2016\)](#) is utilized for prompt filtering, and can also be replaced by any other graph completion
546 methods. We encode the input sequence with a two-layer BERT model, and utilize the sum operation
547 to aggregate the encoded sequence. Finally, a two-layer MLP is applied as the prediction head for the
548 HGC task. During training, we employ the AdamW optimizer and use binary cross-entropy as the
549 loss function. The learnable parameters of the pre-trained DistMult are initialized randomly, while
550 BERT is loaded with pretrained weight parameters. The training process is conducted on an NVIDIA
551 RTX 3090 GPU with 24 GB of memory.

552 **E Related Work**

553 For the task of graph completion/link prediction, methods that learn both the entity and relation
554 representations are categorized into embedding-based and transformer-based, depending on their
555 intrinsic modeling structures.

556 **Embedding-based Methods.** Knowledge graph embedding (KGE) methods employ geometric opera-
557 tions in the vector space to capture the underlying semantics of the graph, such as translation [Bordes
558 et al. \(2013\)](#), bilinear transformation [Yang et al. \(2015\)](#), rotation [Sun et al. \(2018\)](#). Other methods
559 design embeddings from different perspectives. For instance, CompLEX [Trouillon et al. \(2016\)](#)
560 leverages compositionality to model the complex relationships between entities. ConvE [Dettmers
561 et al. \(2018\)](#) utilizes multi-layer convolutional networks on the 2D grid abstracted from the knowledge
562 graph to encode local dependencies. Although conceptually straightforward, these methods encode
563 each entity and relation’s embedded information through a simple vector. The inherent simplicity
564 of embedding-based methods can present challenges in scenarios involving complex reasoning and
565 scarcity of information.

566 **Transformer-based Methods.** Taking account of the relatively weak expression power of the
567 embedding-based methods, several recent works utilize transformers for additional enhanced con-
568 textual information encoding. Some works take the triple as the input and perform tasks such as
569 triple classification and link prediction. For example, KG-BERT [Yao et al. \(2019\)](#) treats triples as
570 textual sequences to inject semantic information and exploits pretrained BERT to learn context-aware
571 embeddings. PKGC [Lv et al. \(2022\)](#) leverages the entity’s semantic information and converts them
572 into natural prompt sentences to address the closed-world assumption (CWA) and incoherent issue.
573 However, the above methods require the scoring of all possible triples in inference, therefore introduc-
574 ing some unnecessary calculation overheads. On the other hand, some other works are designed to
575 directly output the candidate entities. For example, StAR [Wang et al. \(2021\)](#) designs a structure-aware
576 and structure-augmented framework for efficient KGC inference. HittER [Chen et al. \(2021\)](#) extracts
577 context neighbors for the source entity and introduces the additional masked entity prediction task
578 for balanced contextualization. GenKGC [Xie et al. \(2022\)](#) introduces relation-aware demonstration
579 and entity-aware hierarchical decoding for better representation learning. Despite the progress made
580 so far, we notice some implementation gaps in applying the above methods to a knowledge graph
581 and a heterogeneous graph: First, entities in a knowledge graph naturally entitle semantic information,
582 while this is not always true for a heterogeneous graph; Second, the above methods left out the
583 entity/node type information provided in a heterogeneous graph, therefore leaving considerable space
584 for performance improvement. In contrast, our model is designed to not only straightly output the
585 candidate entities, which eliminates the calculation overhead, but also fully utilize the entity and
586 relation type information for better prompting.