# Distributional Inclusion Hypothesis and Quantifications: Probing for Hypernymy in Functional Distributional Semantics

Anonymous ACL submission

#### Abstract

Functional Distributional Semantics (FDS) models the meaning of words by truthconditional functions. This provides a natural representation for hypernymy but no guarantee that it can be learnt when FDS models are trained on a corpus. In this paper, we probe into FDS models and study the representations learnt, drawing connections between quantifications, the Distributional Inclusion Hypothesis (DIH), and the variational-autoencoding objective of FDS model training. Using synthetic data sets, we reveal that FDS models learn hypernymy on a restricted class of corpus that strictly follows the DIH. We further introduce a training objective that both enables hypernymy learning under the reverse of the DIH and improves hypernymy detection from real corpora.

## 1 Introduction

001

006

009

011

013

017

021

033

037

Functional Distributional Semantics (FDS: Emerson and Copestake, 2016; Emerson, 2018) suggests that the meaning of a word can be modelled as a truth-conditional function, whose parameters can be learnt using the distributional information in a corpus (Emerson, 2020a; Lo et al., 2023). Aligning with truth-conditional semantics, functional representations of words are logically more rigorous than vectors (e.g., Mikolov et al., 2013; Pennington et al., 2014; Levy and Goldberg, 2014; Czarnowska et al., 2019) and distributions (e.g., Vilnis and Mc-Callum, 2015, Bražinskas et al., 2018) as concepts are separated from their referents (for a discussion, see: Emerson, 2020b, 2023). On top of its theoretical favour, Lo et al. (2023) also demonstrate FDS models in action and show that they are very competitive in the semantic tasks of semantic composition and verb disambiguation.

Hypernymy is also known as lexical entailment.It is formally defined as the subsumption of extensions between two word senses, which can be modelled with truth-conditional functions. Although

FDS provides the basis for embedding hypernymy, it is not obvious whether hypernymy can be learnt by training an FDS model on a corpus, and if so, what kind of corpus it can successfully learn from. 041

042

043

044

045

046

047

049

051

055

056

057

058

060

061

062

063

064

065

066

067

068

069

070

071

072

073

074

075

076

077

078

079

To acquire hypernymy automatically from a corpus, one way is through the use of distributional information. In this class of methods, hypernymy is learnt in an unsupervised manner given certain hypotheses about the distributional properties of the corpus. One such hypothesis is the Distributional Inclusion Hypothesis (DIH; Weeds et al., 2004; Geffet and Dagan, 2005), which relates lexical entailment of words to the subsumption of the typical contexts they appear with in a corpus.

In §2, we first highlight that while existential quantifications support the DIH, universal quantifications reverse it. In §3, we discuss how FDS can embed hypernymy, formulate our hypothesis that FDS learns hypernymy under the DIH, and introduce a training objective that handles simple universal quantifications. In §4, using synthetic data sets, we show that FDS learns hypernymy under the DIH, and under the reverse of DIH when the new objective is used. In §5, we show that the new objective encodes word generality and improves hypernymy detection by FDS on real corpora.

## 2 Distributional Inclusion Hypothesis and Quantifications

The Distributional Inclusion Hypothesis (DIH) asserts that the typical characteristic features (contexts) of  $r_h$  are expected to appear with  $r_H$  if and only if  $r_H$  is a hypernym of  $r_h$ . Although Geffet and Dagan (2005) report that the DIH is largely valid on a real corpus, it is not deemed fully correct in general as a hyponym can appear in exclusive contexts due to collocational (Rimell, 2014) and pragmatic reasons (Pannitto et al., 2018), and feature inclusion has been found to be selective (Roller et al., 2014). In this section, we describe how quantifications can also be pivotal to the hypothesis.

$$animal \{ \stackrel{\text{ARG1}}{\leftarrow} grow \}$$

$$animal \{ \stackrel{\text{ARG1}}{\leftarrow} furry \}$$

$$dog \{ \stackrel{\text{ARG1}}{\leftarrow} bark \} bat \{ \stackrel{\text{ARG1}}{\leftarrow} fly \}$$

Figure 1: A taxonomic hierarchy of nouns. Next to each noun is the set of contexts that are applicable to the extension of it and those of its descendants (e.g., all dogs are furry, but not all animals.).

Corpus 1 (DIH)	Corpus 2 (rDIH)		
a dog barks	every dog barks		
a mammal barks	every dog is furry		
an animal barks	every dog grows		
a bat flies	every bat flies		
a mammal flies	every bat is furry		
an animal flies	every bat grows		
a mammal is furry	every mammal is furry		
an animal is furry	every mammal grows		
an animal grows	every animal grows		

Table 1: Corpora generated from the hierarchy in Fig. 1. Existential and universal quantifications result in two corpora that follow the DIH and rDIH respectively.

087

089

094

096

100

101

102

104

105

106

107

While Geffet and Dagan (2005) consider syntaxbased context, we suggest that contexts based on semantic representation are more suitable since syntactic differences do not necessarily contribute to semantic ones (e.g., passivizations and inversions), and the subject of concern should be semantics. We use Dependency Minimal Recursion Semantics (DMRS; Copestake et al., 2005; Copestake, 2009) as the semantic representation, which is derived using the English Resource Grammar (ERG; Flickinger, 2000, 2011). Fig. 2 shows the predicate– argument structure of an example DMRS graph. If  $r_i \leftarrow \frac{ARG[a]}{c_j} r_j$  exists in the DMRS graph of a sentence in the corpus, we can say that  $r_i$  appears in the context  $\leftarrow \frac{ARG[a]}{c_j} r_j$ .

Consider a corpus as a partial description of a world. Distributional properties would depend on how the world is described. Here, we consider a corpus of simple sentences in the form '*[quantifier] [noun] [context word]*'. Take the taxonomic hierarchy in Fig. 1 as an example, where each noun has a set of applicable contexts. If we want to generate existentially quantified statements that are true, then: (1) a noun can appear in its hypernyms' contexts, e.g., '*a dog grows*', where  $\leftarrow^{ARG1}$  grow is applicable to *animal*; and (2) a noun can appear in its hyponyms' contexts, e.g., '*an animal barks*', where  $\stackrel{\text{ARG1}}{\longleftarrow}$  bark is applicable to dog. If we only generate (2) and restrict (1) so that contexts that are broadly applicable are not used with more specific nouns, this creates a corpus that follows the DIH. Corpus 1 of Table 1 shows an example.<sup>1</sup>

108

109

110

111

112

113

114

115

116

117

118

119

120

121

122

123

124

125

126

127

128

129

130

131

133

134

135

136

137

138

139

140

141

142

143

144

145

146

147

148

149

150

151

152

In contrast, generating sentences with universal quantifications results in a corpus that follows the *reverse* of the DIH (rDIH), as in Corpus 2, where the set of contexts of *mammal* is a subset of that of *dog*. Consequently, methods that rely on the DIH as a cue for hypernymy would be undermined.

In §5, we use these processes to generate corpora which strictly align with the DIH or rDIH. Corpora with more complex sentence structures would require a richer world model than can be encoded in a taxonomic hierarchy like Fig. 1. For instance, with a restricted relative clause, 'every dog that is trained is gentle' does not entail 'every Chihuahua is gentle' even if Chihuahua is a hyponym of dog, as the universal quantifier applies only to trained dogs. We also disregard negations because they can co-occur nearly freely, effectively making a context word in the negated scope uninformative. For example, 'a dog does not \_\_\_\_' is much less selective than 'a dog \_\_\_\_'.

## 3 Functional Distributional Semantics

In this section, we introduce Functional Distributional Semantics (FDS), discuss hypernymy representation in FDS and explain how FDS can be adapted to handle quantifications. We follow Lo et al. (2023)'s FDS implementation which is briefly described here.

#### 3.1 Model-Theoretic Semantics

FDS is motivated by model-theoretic semantics, which sees meaning in terms of an extensional model structure that consists of a set of *entities*, and a set of *predicates*, each of which is true or false of the entities. In parallel, FDS represents an entity by a *pixie* which is taken to be a high-dimensional feature vector, and represents a predicate by a truthconditional *semantic function* which takes pixie(s) as input and returns the probability of truth.

#### 3.2 Probabilistic Graphical Models

The framework is formalized in terms of a family of probabilistic graphical models. Each of them

<sup>&</sup>lt;sup>1</sup>Without the restriction on (1), exhaustively generating true assertions generates a corpus where the DIH does not hold between nouns in a unary chain (e.g., *animal* and *mammal* in Fig. 1), which would appear in the same set of contexts.



Figure 2: Probabilistic graphical model of FDS for generating words in an SVO triple '*postman deliver mail*'. Only  $R_1 = postman$ ,  $R_2 = deliver$ , and  $R_3 = mail$  are observed.

describes the generative process of predicates in 153 the semantic graph of a sentence. Fig. 2 illustrates 154 the process of generating the words given the ar-155 gument structure  $R_1 \xleftarrow{\operatorname{ARG1}} R_2 \xrightarrow{\operatorname{ARG2}} R_3$ . First, a 156 pixie  $Z_i \in \mathbb{R}^d$  is generated for each node in the graph, together representing the entities described 158 by the sentence. Then, for each pixie  $Z_i$ , a truth 159 value  $T_{Z_j}^{(r_i,0)}$  is generated for each predicate  $r_i$  in the vocabulary  $\mathcal{V}$ ; and for each pair of nodes connected as  $R_j \xrightarrow{\text{ARG}a} R_k$  whose corresponding pix-161 162 ies are  $Z_j$  and  $Z_k$ , a truth value  $T_{Z_j,Z_k}^{(r_i,a)}$  is generated for each predicate  $r_i$  in the vocabulary. Finally, a 164 single predicate  $R_i$  is generated for each pixie  $Z_i$ 165 conditioned on the truth values. 166

#### 3.3 Semantic Functions

As mentioned in §3.1, instead of treating a predicate as an indicator function, FDS models the probability that it is true of the pixie(s) with unary (in (1)) and binary semantic functions (in (2)). This allows the model to account for vagueness.

$$P\left(T_{Z_j}^{(r_i,0)} = \top \middle| z_j\right) = t^{(r_i,0)}(z_j) \tag{1}$$

$$P\left(T_{Z_j,Z_k}^{(r_i,a)} = \top \middle| z_j, z_k\right) = t^{(r_i,a)}(z_j, z_k) \quad (2)$$

The functions are implemented as linear classifiers:

$$t^{(r_i,0)}(z_j) = S\left(v^{(r_i,0)^{\top}} z_j + b^{(r_i,0)}\right) \quad (3)$$

$$t^{(r_i,a)}(z_j,z_k) =$$

177 
$$S\left(v_1^{(r_i,a)^{\top}} z_j + v_2^{(r_i,a)^{\top}} z_k + b^{(r_i,a)}\right)$$
(4)

178

168

169

172

173

174

175

176

where S denotes the sigmoid function.

#### 3.4 Representing Hypernymy

In truth-conditional semantics, for a set of entities D,  $r_H$  is a hypernym of  $r_h$  if and only if

$$\forall x \in D \colon r_h(x) \Longrightarrow r_H(x) \tag{5}$$

179

180

181

182

183

184

186

187

188

190

191

192

193

194

196

197

198

199

200

201

203

204

207

209

210

211

212

213

214

215

216

217

218

220

221

Although FDS provides truth-conditional interpretations of words, it is not straightforward to define hypernymy in FDS where predicates are probabilistic and work over high-dimensional pixies. One way is to translate (5) to probabilistic counterpart for a score on hypernymy  $P\left(T_Z^{(r_H,0)}=\top \mid T_Z^{(r_h,0)}=\top\right)$ . However, this conditional probability is unavailable since only  $P\left(T_Z^{(r_H,0)}=\top \mid z\right)$  and  $P\left(T_Z^{(r_h,0)}=\top \mid z\right)$  are modelled by FDS.

Another way is to interpret the probability model from a fuzzy set perspective and use fuzzy set containment (Zadeh, 1965) for representing hypernymy:

$$\forall z : t^{(r_H,0)}(z) > t^{(r_h,0)}(z) \tag{6}$$

Note that if we consider all  $z \in \mathbb{R}^d$ , (6) can only be true when  $v^{(r_h,0)} = kv^{(r_H,0)}$  where  $k \neq 0$ , which is impossible to be obtained in practice from model training. Therefore, we restrict the pixie space and only consider pixies in a unit hypersphere or hypercube to be meaningful. With (3) and (4),  $r_H$ is considered the hypernym of  $r_h$  if and only if  $s(r_h, r_H) > 0$  in (7), where  $p \in \{1, 2\}$  (derivation in Appendix A). Cheng et al. (2023) also use this score for hypernymy.

$$s(r_h, r_H) = b^{(r_H, 0)} - b^{(r_h, 0)} - \left\| v^{(r_H, 0)} - v^{(r_h, 0)} \right\|_p$$
(7)

Note that the transitivity of (5) is paralleled (derivation in Appendix B):

$$s(r_1, r_2) > 0 \land s(r_2, r_3) > 0 \implies s(r_1, r_3) > 0$$
(8)

Having hypernymy representation built into a distributional model allows generalization out of missing information. For example, when the (r)DIH does not hold between *dog* and *mammal* in a corpus, knowing that both *dog* and *fox* share the same contexts in a corpus (e.g.,  $\{\langle ARG1 \ bark\}\}$ ) is indicative that they share common hypernyms, e.g., *mammal*. This also applies to hyponymy, e.g., *machine* and *system* sharing  $\{\langle ARG1 \ complex\}$  and sharing *computer* as their hyponym. Such generalization power is largely absent in models based on strict contexts subsumption.

#### 3.5 **Original Training Objective**

FDS models are trained using the variationalautoencoding method on simplified DMRS graphs where quantifiers and scopal information are removed from the graphs before training, leaving us with just the predicate-argument structure. The approximate posterior distribution of pixies  $q_{\phi}$  is taken to be n spherical Gaussian distributions, each with mean  $\mu_{Z_i}$  and covariance  $\sigma_{Z_i}^2 I$ . Given an observed DMRS graph G with n pixies  $Z_1 \dots Z_n$ , we maximize (9), reformulated from the  $\beta$ -VAE (Higgins et al., 2017).

> $\mathcal{L} = \sum_{i=1}^{n} \mathcal{C}_{i} + \sum_{r_{i} \xrightarrow{\operatorname{ARG}[a]} r_{j} \operatorname{in} G} \mathcal{C}_{i,j,a}$ (9)  $-\frac{d}{2}\sum_{i=1}^{n}\beta_{1}\mu_{Z_{i}}^{2}+\beta_{2}\left(\sigma_{Z_{i}}^{2}-\ln\sigma_{Z_{i}}^{2}\right)$

The last term in (9) is the regularization term for the approximate posterior, and the first two terms aim to maximize the truthness of observed predicates and the falsehood of K negatively sampled ones  $r'_k$ over the inferred pixie distribution  $q_{\phi}$ , by

$$C_{i} = \ln \mathbb{E}_{q_{\phi}} \left[ t^{(r_{i},0)}(z_{i}) \right] + \sum_{k=1}^{K} \ln \mathbb{E}_{q_{\phi}} \left[ 1 - t^{(r'_{k},0)}(z_{i}) \right]$$

$$C_{i,j,a} = \ln \mathbb{E}_{q_{\phi}} \left[ t^{(r_{i},a)}(z_{i},z_{j}) \right] + \sum_{l=1}^{K} \ln \mathbb{E}_{q_{\phi}} \left[ 1 - t^{(r'_{k},a)}(z_{i},z_{j}) \right]$$
(11)
(11)

Both the local predicate–argument structure of  
each predicate and global topical information in  
the graph are used for variational inference. For in-  
stance, the approximate posterior distribution of the  
pixie 
$$Z_1$$
 of *postman* in Fig. 2 is inferred from the  
direct argument information,  $\langle ARG1 \rangle$  *deliver*, and  
the indirect topical predicate,  $\not\leftarrow$  mail.

**Our Hypothesis.** If the training corpus strictly follows the DIH, hypernymy can be learnt by FDS models. The intuition behind our hypothesis is elaborated in Appendix C.

#### **Proposed Objective for Universal** 3.6 Quantifications

FDS assumes that each observed predicate refers to only one point in the pixie space and offers no tools for dealing with regions. We propose a method to allow optimizations of semantic functions with respect to a region in the pixie space, thus enabling FDS to handle simple sentences with universal quantifications. Essentially, we add the following  $\forall$ -objective to the original objective in **(9**):

259

260

261

262

263

264

265

267

268

271

272

274

275

276

277

278

279

282

283

285

286

287

288

291

292

293

$$\mathcal{L}_{\forall} = \sum_{r_j \xleftarrow{}^{\text{ARG}[a]} r_i \text{ in } G} s_a(r_i, r_j) + \mathcal{U}_{i,j,a} \quad (12)$$

where  $r_i$  is a predicate whose referent is universally quantified, and

$$s_{a}(r_{i}, r_{j}) = b^{(r_{i}, a)} - b^{(r_{j}, 0)} - \left\| v_{2}^{(r_{i}, a)} - v^{(r_{j}, 0)} \right\|_{p}$$
(13) 269

$$\mathcal{U}_{i,j,a} = \sum_{k=1}^{K} \min\left(0, -s_0(r_i, r'_k)\right) + \sum_{k=1}^{K} \min\left(0, -s_{a''_k}(r''_k, r_j)\right)$$
(14) 2

Note that (13) is modified based on (7), previously defined for classifying hypernymy.

To explain (12), consider the sentence '*every* dog barks' as an example. The first term inside the summation in (12) enforces that extension of  $r_i$  is a subset of that of prototypical argument a of  $r_i$ , i.e., the set of dogs should be contained in the set of agents that barks. The second term, described in (14), incorporates K randomly generated negative samples.  $r'_k$  is a noun, which is a negative sample for  $r_j$ .  $r''_k$  is a verb or adjective and  $a''_k$  is an argument role, together form a negative sample for  $r_i$  and a. Then, (14) requires that it is false to universally quantify the referents of the noun  $r'_k$  in  $r'_k \xleftarrow{\operatorname{ARG}[a]} r_i$  and  $r_j$  in  $r_j \xleftarrow{\operatorname{ARG}[a''_k]} r''_k$ . For the example, both of the following sentences are considered false: 'every dog is owned' and 'every cat barks', where  $r'_k = cat$ ,  $r''_k = own$  and  $a''_k = 2$ .

#### 4 **Experiments on Synthetic Data Sets**

Testing our hypothesis in §3.5 and the effectiveness of the new objective for universal quantifications in §3.6 requires corpora that strictly follow the DIH or rDIH, which is impractical for real corpora. Therefore, we create a collection of synthetic data sets and perform experiments on them.

236

224

225

227

228



240

241

245

246

247

248

252



Figure 3: Examples of the topologies of the synthetic taxonomic hierarchies.

#### 4.1 Synthetic Data Sets under the (r)DIH

297

301

307

310

311

312

313

316

319

320

323

327

329

331

332

333

Each of the synthetic data sets consists of a taxonomic hierarchy of nouns and a corpus, created using the following procedure:

- 1. **Create a taxonomic hierarchy.** Define a set of nouns, the hypernymy relations of them, and the contexts applicable to its extension and those of its hyponyms (as in Fig. 1).
- 2. Choose a hypothesis. The DIH or rDIH.
- 3. **Create a corpus.** Create sentences in the form *`[quantifier] [noun] [context word]*' following the chosen hypothesis and the defined hierarchy (as in Table 1).

### 4.1.1 Topology of Hierarchy

Different topologies of hierarchy lead to different distributional usage of words, thus possibly varying representations learnt for hypernymy. For example, a noun can have multiple hypernyms (e.g., *dog* is the hyponym of both *pet* and *mammal*), or share overlapping contexts with another noun far in the hierarchy (e.g., both *bat* and *airplane*  $\stackrel{\text{ARG1}}{\leftarrow}$  *fly*).

To test the robustness of FDS models for learning hypernymy, we experiment with a range of topologies. Fig. 3 exemplifies the five classes of topologies used. We expect that directed acyclic graphs  $(H_{\text{DAG}} \text{ and } H'_{\text{DAG}})$  be harder topologies than trees  $(H_{\text{tree}} \text{ and } H'_{\text{tree}})$ , and topologies with overlapping contexts  $(H'_{\text{tree}} \text{ and } H'_{\text{DAG}})$  be harder than those without  $(H_{\text{tree}} \text{ and } H'_{\text{DAG}})$ . In addition, we test  $H_{\text{chains}}$  with pixie dimensionality d = 2. A 2-D pixie space allows lossless visualization of the semantic functions. To test hypernymy learning at scale on an actual hierarchy, we make use of Word-Net (Miller, 1995; Fellbaum, 1998) and test our models on the WordNet's hierarchy  $(H_{\text{WN}})$ .

Every node in the hierarchy consists of a noun and a semantic context. The topology of the  $H_{chains}$ used in the experiment is exactly as depicted in Fig. 3.  $H_{WN}$  is created out of the synset *entity.n.01* in WordNet, which is the root, and its hyponymic synsets. This results in 74,374 nodes with 663,492 hypernymy pairs. We randomly sample 663,492 pairs from the remaining pairs as negative instances for evaluation. For the remaining hierarchies, each of them consists of 153 nodes with a height of 5. For  $H_{\text{tree}}$ , the first level is a root node, and a node at the  $h^{\text{th}}$  level has (h + 1) direct children.  $H_{\text{tree'}}$  is created from  $H_{\text{tree}}$  by choosing 5 pairs of nodes and making each pair share a context set.  $H_{\text{DAG}}$  and  $H'_{\text{DAG}}$  are created from  $H_{\text{tree}}$  and  $H_{\text{tree'}}$ respectively by choosing 5 pairs of nodes, where the nodes of each pair are at different levels, and make the higher level node the direct parent of the lower level one. 336

337

338

339

340

341

342

343

344

345

346

349

350

351

352

353

354

357

358

359

360

361

362

363

364

365

366

367

368

369

370

371

372

373

374

375

#### 4.2 FDS Models Training

We experiment with two variations of FDS training: FDS is trained using the original objective in (9) whereas FDS $\forall$  incorporates the  $\forall$ -objective following §3.6. The hypernymy score of each model, given by (7), is averaged over two runs of different random seed. We empirically find that setting p = 1 in (13) and p = 2 in (7) almost always give the best performances, and we only report the results in this setup. Other than the newly introduced training objective, training of the models largely follows that of Lo et al. (2023). No hyperparameter search is conducted due to the large number of experiments (details described in Appendix D).

#### 4.3 Evaluation on Hypernymy Detection

We test if a model trained on the corpus learns to identify hypernymy defined in the hierarchy that generates the corpus. Concretely, a model is asked to give a score of hypernymy between every pair of nouns using (7). Performance is then measured by the area under the receiver operating characteristic curves (AUC). Unlike average precision, AUC values do not reflect changes in the distribution of classes, which is favourable since we are comparing models' performances across varying class distributions generated from different topologies.

378

379

38

- 501

384

387

390

391

393

397

400

401

402

403

404

405

invCL
$$(r_1, r_2) = \sqrt{\text{CL}(r_1, r_2)(1 - \text{CL}(r_2, r_1))}$$
  
where CL $(r_1, r_2) = \frac{\sum_i \min\left(u_i^{(r_1)}, u_i^{(r_2)}\right)}{\sum_i u_i^{(r_1)}}$ 

invCL (Lenci and Benotto, 2012):

They measure the context inclusion of  $r_1$  by  $r_2$  and invCL measures also the non-inclusion of  $r_2$  by  $r_1$ . Their distributional space is constructed by first counting co-occurrences of adjacent predicates in the preprocessed DMRS graphs, then transforming the resulting matrix using positive pointwise mutual information. Each row vector  $u^{(r_i)}$  in the transformed matrix represents a predicate  $r_i$ .

We include two distributional methods for hy-

pernymy detection based on the DIH in the experi-

ments, namely WeedsPrec (Weeds et al., 2004) and

 $\text{WeedsPrec}(r_1, r_2) = \frac{\sum_i u_i^{(r_1)} \mathbb{1}_{u_i^{(r_2)} > 0}}{\sum_i u_i^{(r_1)}}$ 

Model	$H_{\rm chains}$	$H_{\text{tree}}$	$H_{\rm tree}'$	$H_{\rm DAG}$	$H_{\rm DAG}^\prime$	$H_{\rm WN}$
Fds	.990	.994	.995	.995	.995	.940
$FDS \forall$	.925	.206	.210	.214	.221	.788
WeedsPrec	1.000	1.000	1.000	1.000	1.000	1.000
invCL	1.000	1.000	1.000	1.000	.999	1.000

Table 2: AUC of models trained on the DIH corpora.

Model	$H_{\text{chains}}$	$H_{\text{tree}}$	$H_{\rm tree}'$	$H_{\rm DAG}$	$H_{\rm DAG}^\prime$	$H_{\rm WN}$
FDS	.876	.842	.793	.752	.688	.444
Fds∀	.988	.983	.978	.981	.977	.675
WeedsPrec	.900	.675	.619	.613	.556	.809
invCL	.900	.355	.280	.236	.276	.564

Table 3: AUC of models trained on the rDIH corpora.

Table 2 and Table 3 show the results of FDS models when trained on the DIH and rDIH corpora respectively (a visualization of results on  $H_{chains}$  is provided in Appendix E). As expected, FDS, WeedsPrec and invCL are shown to work on the DIH corpora, and only FDS $\forall$  works on the rDIH corpora. Reversing the FDS models on respective corpora yields substantially worse performances. In particular, FDS $\forall$  attains AUCs of about 0.2 on the DIH corpora means hypernymy predictions are even mostly reversed, which in turn reflects the effectiveness of the universal objective when FDS $\forall$  interprets the subsumption of contexts reversely based on the rDIH. Moreover, hierarchies with overlapping contexts and multiple direct hypernyms are not harder for FDS than those without. Scaling up to the huge WordNet hierarchy  $H_{WN}$  results in a slight drop in AUC for FDS on the DIH corpus, and markedly worse performances for FDS $\forall$ . While our preset hyperparameters work nicely on all other settings, it is possible that FDS $\forall$  requires a different set of hyperparameters to perform optimally on the rDIH corpora generated from huge hierarchies.

406

407

408

409

410

411

412

413

414

415

416

417

418

419

420

421

422

423

424

425

426

427

428

429

430

431

432

433

434

435

436

437

438

439

440

441

442

443

444

445

446

447

Model	$H_{\text{chains}}$	$H_{\text{tree}}$	$H_{\rm DAG}$	$H_{\rm WN}$
Fds	.947	.991	.990	.946
$FDS_{\forall}$	.956	.358	.378	.183
FDS <sub>(\()</sub>	.999	.986	.984	.992
WeedsPrec	.950	.996	.993	.993
invCL	.958	.891	.864	.868

Table 4: AUC of models trained on combined corpora.

We further produce new corpora by combining the DIH and rDIH corpus of each topology. The resulting corpora still follow the DIH. In this setup, instead of applying the same FDS training objective across the whole corpus, the  $\forall$ -objective can be added only when there is a universal quantifier (e.g., *every*). We name this FDS model  $FDS_{(\forall)}$ . Table 4 shows the results. The DIH methods (FDS, WeedsPrec, and invCL) perform well as expected, whereas invCL performs worse since it also measures non-inclusion which is undermined by the rDIH half. There are several interesting insights into FDS models. First, FDS still works on the corpora with an rDIH half because it is still valid to say 'a dog grows', as mentioned in §2. Second,  $FDS_{(\forall)}$  is as good as FDS across topologies and even better on  $H_{\text{chains}}$  and  $H_{\text{WN}}$ . This implies that the  $\forall$ -objective indeed captures simple universal quantifications and can be used compatibly with the original training method on a corpus with varying quantifications. Third, on  $H_{WN}$ ,  $FDS_{(\forall)}$ performs better than FDS on the DIH corpora and much better than  $FDS_{\forall}$  on the rDIH corpora. This reflects that the ∀-objective is more effective when mixed with the original mode of training.

## 4.4 Evaluation on Distributional Generalization

We also test if the distributional generalization power mentioned in §3.4 exists in FDS. We construct a new corpus from a hierarchy with removed hypernymy information. Fig. 4 illustrates the idea with an example hierarchy of nouns and the contexts that would appear with the nouns in the new corpora obtained. If upward (downward) distri-



Hyp.	Noun	Contexts
DIH	$r_4 r_5$	${c_5, c_7} {c_5, c_8}$
rDIH	$r_4 r_5$	$ \begin{cases} c_1, c_3, c_5 \\ \{c_1, c_2, c_5 \end{cases} $

(a) An example taxonomy with missing relations. The siblings  $r_4$  and  $r_5$  have the same context  $c_5$ . Dashed lines show the relations to be removed from training.

449

450

451

452

453

454

455

456

457

458

459

460

461

462

463

464

465

466

467

468

469

470

471

(b) Contexts that  $r_4$  and  $r_5$  in a would appear with in the new (r)DIH corpora.

Figure 4: Illustration of the setup for testing distributional generalization.

butional generalization exists in a model, based on that  $r_4$  and  $r_5$  share  $c_5$  as their contexts, it should identify the hypernyms (hyponyms) of  $r_5$ ( $r_4$ ) as the *candidate hypernyms* (hyponyms) of  $r_4$  ( $r_5$ ) after training on the new corpus. That is, we expect  $\forall r_j \in \{r_5, r_6, r_7, r_8\}: s(r_4, r_2) >$  $s(r_4, r_j)$  if upward generalization exists, and  $\forall r_j \in$  $\{r_1, r_2, r_3, r_4, r_6\}: s(r_7, r_5) > s(r_j, r_5)$  if downward exists in FDS.

In our experiments, we sample five nouns from the  $H'_{\text{DAG}}$  hierarchy. Then, for each of these nouns  $\tilde{r}$ , we equate the contexts set of  $\tilde{r}$  to that of one of its siblings and remove the hypernymy (hyponymy) information of their common parent (daughter) from  $\tilde{r}$  when creating the new corpora.

Model	Hypothesis	Upward	Downward
Fds	DIH	.922	.742
$F D s_\forall$	rDIH	.976	.998

Table 5: Mean AUC for distributional generalizations.

For each  $\tilde{r}$ , we compute the hypernymy score of between  $\tilde{r}$  and each of the candidate hypernyms, and between  $\tilde{r}$  and a random noun. We measure the performance with mean AUC, averaged over the five chosen  $\tilde{r}$ . Table 5 shows that both upward and downward distributional generalizations exist when the corpus follows either the DIH or rDIH, and to a larger extent on the rDIH corpus.

#### 4.5 Summary

The experimental results confirm that: (1) the original FDS models learn hypernymy under the DIH,
(2) the proposed ∀-objective captures universal quantifications and enables hypernymy learning under the rDIH, and (3) FDS models can generalize about nouns with incomplete contexts in a

corpus using distributional information.

#### 5 Experiments on Real Data Sets

Seeing how FDS performs on restricted synthetic data sets is helpful for understanding models' behaviour but it does not immediately tell us more about hypernymy learning from open classes of sentences. Therefore, we perform further experiments using a real corpus and data sets for hypernymy. 478

479

480

481

482

483

484

485

486

487

488

489

490

491

492

493

494

495

496

497

498

499

500

501

502

503

504

505

506

507

508

509

510

511

512

513

514

515

516

517

518

519

520

521

522

523

524

#### 5.1 FDS Models Training

**Training Data.** FDS models are trained on Wikiwoods (Flickinger et al., 2010; Solberg, 2012), which provide linguistic analyses of 55m sentences (900m tokens) in English Wikipedia. Each of the sentences was parsed by the PET parser (Callmeier, 2001; Toutanova et al., 2005) using the 1212 version of the ERG, and the parses are ranked by a ranking model trained on WeScience (Ytrestøl et al., 2009). We extract the DMRS graphs from Wikiwoods using Pydelphin<sup>2</sup> (Copestake et al., 2016). After preprocessing, there are 36m sentences with 254m tokens.

**Model Configurations.** Although quantifications are annotated in Wikiwoods, neither of the proposed training objectives is entirely applicable in general. For example, even for a sentence of modest complexity like '*every excited dog barks*', it requires a universal quantification over the intersection of the set of dogs and excited entities. However, set intersection is not modelled by FDS. In our experiments, we apply either FDS or FDS<sub> $\forall$ </sub> described in §4.2 to every training instance. We also test an additional model FDS<sub> $\forall/2$ </sub> where the  $\forall$ objective is scaled by 0.5. Each model is trained for 1 epoch and the results of each model are averaged over two random seeds as discussed in §4.2.

#### 5.2 Evaluation Method

We test the trained models on four English hypernymy data sets for nouns, namely Kotlerman2010 (Kotlerman et al., 2010), LEDS (Baroni et al., 2012), WBLESS (Weeds et al., 2014), and EVA-Lution (Santus et al., 2015). Each of them consists of a set of word pairs, each with a label indicating whether the second word is a hypernym of the first word. We removed the out-of-vocabulary instances from all data sets, and non-nouns from EVALution during the evaluation. Table 6 reports the statistics of the test sets data. We report the AUC as in §4.

<sup>&</sup>lt;sup>2</sup>https://github.com/delph-in/pydelphin

Test Set	# Positive	# Negative
Kotlerman2010	880 [831]	2058 [1919]
LEDS	1385 [1344]	1385 [1342]
WBLESS	834 [830]	834 [813]
Evalution	1592 [1352]	4561 [3241]

Table 6: Class distributions of test sets. In brackets are the numbers after removal of OOV instances and non-nouns.

In addition, we use WBLESS for further performance analysis, which provides categorizations of the negative instances. Each of the negative instances is either a hyponymy pair, co-hyponymy pair, meronymy pair, or random pair.

## 5.3 Baselines

525

526

527

528

529

530

532

533

535 536

537

538

539

540

541

542

543

545

546

547

Following Roller et al. (2018), we implement five distributional methods and train them on Wiki-woods using the distributional space described in §4.3. Apart from the two DIH measures in §4.3, we use SLQS (Santus et al., 2014), a word generality measure that rests on another hypothesis that general words mostly appear in uninformative contexts:

$$SLQS(r_1, r_2) = 1 - \frac{E_{r_1}}{E_{r_2}}$$
  
where  $E_{r_i} = \text{median}_{i=1}^N [H(c_j)]$ 

For each word  $r_i$ , the median of the entropies of N most associated contexts (as measured by local mutual information) is computed, where  $H(c_j)$  denotes the Shannon entropy of the associated context  $c_j$ . Then, SLQS compares the generality of two words by their medians. N is chosen to be 50 following Santus et al. (2014). We also include cosine similarity (Cosine) of  $u^{(r_1)}$  and  $u^{(r_2)}$ , and SLQS–Cos, which multiplies the SLQS measure by Cosine since the SLQS measure only considers generality but not similarity.

#### 5.4 Results

Table 7 shows the results on the four test data sets. 552 The DIH baselines are competitive and nearly out-553 perform all models across the test sets.  $FDS_{\forall}$  and 554  $FDS_{\forall/2}$  both outperform FDS considerably across the test sets. This reflects that including the pro-556 posed  $\forall$ -objective in training is useful for extracting hypernymy information in a corpus. Compared to 558 the 2.7-billion-token corpus used by Santus et al. 559 (2014) in training SLQS, we suggest that the Wiki-560 woods corpus is too small for SLQS to obtain meaningful contexts of the median entropy: setting N

Model	Kotlerman2010	LEDS	WBLESS	Evalution
Cosine	.701	.782	.620	.526
WeedsPrec	.674	.897	.709	.650
invCL	.679	.905	.707	.620
SLQS	.491	.480	.568	.532
SLQS-Cos	.489	.477	.557	.532
FDS	.473	.650	.508	.459
FDS <sub>∀/2</sub>	.558	.759	.660	.583
$FDS_{\forall}$	.550	.735	.655	.554

Table 7: AUC on the test sets.

to be small results in frequent contexts that are not representative of the nouns, whilst setting it large would require a disproportionate number of contexts for the infrequent words.

Model	Hyponymy	Co-hyponymy	Meronymy	Random
Cosine	.511	.369	.683	.924
WeedsPrec	.754	.615	.631	.843
invCL	.745	.568	.652	.872
SLQS	.606	.551	.590	.524
SLQS-Cos	.581	.525	.574	.547
Fds	.596	.288	.561	.587
$FDS_{\forall/2}$	.783	.612	.549	.704
$F D s_\forall$	.783	.625	.527	.691

Table 8: AUC on the sub-categories of WBLESS.

Table 8 shows the results on the WBLESS subcategories. It is shown that  $FDS_{\forall}$  is stronger than the DIH baselines in distinguishing between hyponymy and hypernymy pairs, and between cohyponymy and hypernymy pairs, while weaker for meronymy or random pairs.  $FDS_{\forall}$  and  $FDS_{\forall/2}$  outperform FDS in three out of the four sub-categories, with much higher distinguishing power for cohyponymy and hyponymy. These imply that the  $\forall$ -objective makes FDS more sensitive to the relative generality than the similarity of word pairs.

### 6 Conclusion

We have discussed how Functional Distributional Semantics (FDS) can provide a truth-conditional representation for hypernymy and demonstrate that it is learnable from the distributional information in a corpus. On synthetic data sets, we confirm that FDS learns hypernymy under the Distributional Inclusion Hypothesis (DIH), and under the reverse of the DIH if the proposed objective for universal quantifications is applied. On real data sets, the proposed objective substantially improves FDS performance on hypernymy detection. We hope that this work provides insights into FDS models and hypernymy learning from corpora in general. 565 566

567

568

569

570

571

572

573

574

575

576

577

578

579

580

581

582

583

584

585

586

587

588

589

590

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

643

## Limitations

592

610

611

615

617

618

619

627

631

633

634 635

636

637

638

641

The proposed representation of hypernymy in FDS compares the semantic functions of DMRS pred-594 icate pairs. Following previous implementations 595 of Functional Distributional Semantics, a semantic 596 function is a linear classifier. Consequently, each 598 DMRS predicate is assumed to have only one sense. Modelling polysemy would require more expressive parametrizations of semantic functions, which can pose additional challenges to model training, and the hypernymy representation would possibly need to be revised. Such an approach is considered out of the scope of this work.

## Ethics Statement

We anticipate no ethical issues directly stemming from our experiments.

## References

- Marco Baroni, Raffaella Bernardi, Ngoc-Quynh Do, and Chung-chieh Shan. 2012. Entailment above the word level in distributional semantics. In Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics, pages 23– 32, Avignon, France. Association for Computational Linguistics.
  - Arthur Bražinskas, Serhii Havrylov, and Ivan Titov. 2018. Embedding words as distributions with a Bayesian skip-gram model. In *Proceedings of the* 27th International Conference on Computational Linguistics, pages 1775–1789, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Ulrich Callmeier. 2001. Efficient parsing with largescale unification grammars. Master's thesis, Universität des Saarlandes, Saarbrücken, Germany.
- Xuyou Cheng, Michael Schlichtkrull, and Guy Emerson. 2023. Are embedded potatoes still vegetables? On the limitations of WordNet embeddings for lexical semantics. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 8763–8775, Singapore. Association for Computational Linguistics.
- Ann Copestake. 2009. **Invited Talk:** slacker semantics: Why superficiality, dependency and avoidance of commitment can be the right way to go. In *Proceedings of the 12th Conference of the European Chapter of the ACL (EACL 2009)*, pages 1–9, Athens, Greece. Association for Computational Linguistics.
- Ann Copestake, Guy Emerson, Michael Wayne Goodman, Matic Horvat, Alexander Kuhnle, and Ewa Muszyńska. 2016. Resources for building applications with dependency Minimal Recursion Semantics. In *Proceedings of the Tenth International Conference*

*on Language Resources and Evaluation (LREC'16)*, pages 1240–1247, Portorož, Slovenia. European Language Resources Association (ELRA).

- Ann Copestake, Dan Flickinger, Carl Pollard, and Ivan A. Sag. 2005. Minimal Recursion Semantics: An introduction. *Research on Language and Computation*, 3(2-3):281–332.
- Paula Czarnowska, Guy Emerson, and Ann Copestake. 2019. Words are vectors, dependencies are matrices: Learning word embeddings from dependency graphs. In Proceedings of the 13th International Conference on Computational Semantics - Long Papers, pages 91–102, Gothenburg, Sweden. Association for Computational Linguistics.
- Guy Emerson. 2018. Functional Distributional Semantics: Learning Linguistically Informed Representations from a Precisely Annotated Corpus. Ph.D. thesis, University of Cambridge.
- Guy Emerson. 2020a. Autoencoding pixies: Amortised variational inference with graph convolutions for functional distributional semantics. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3982–3995, Online. Association for Computational Linguistics.
- Guy Emerson. 2020b. What are the goals of distributional semantics? In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7436–7453, Online. Association for Computational Linguistics.
- Guy Emerson. 2023. Probabilistic lexical semantics: From Gaussian embeddings to Bernoulli Fields. In Jean-Philippe Bernardy, Rasmus Blanck, Stergios Chatzikyriakidis, Shalom Lappin, and Aleksandre Maskharashvili, editors, *Probabilistic Approaches to Linguistic Theory*, pages 65–122. CSLI Publications.
- Guy Emerson and Ann Copestake. 2016. Functional distributional semantics. In *Proceedings of the 1st Workshop on Representation Learning for NLP*, pages 40– 52, Berlin, Germany. Association for Computational Linguistics.
- Christiane Fellbaum, editor. 1998. *WordNet: An Electronic Lexical Database* [Website]. Massachusetts Institute of Technology (MIT) Press.
- Dan Flickinger. 2000. On building a more efficient grammar by exploiting types. *Natural Language Engineering*, 6(1):15–28.
- Dan Flickinger. 2011. Accuracy vs. robustness in grammar engineering. In Emily M. Bender and Jennifer E. Arnold, editors, *Language from a Cognitive Perspective: Grammar, Usage, and Processing*, chapter 3, pages 31–50. Center for the Study of Language and Information (CSLI) Publications.
- Dan Flickinger, Stephan Oepen, and Gisle Ytrestøl. 2010. WikiWoods: Syntacto-semantic annotation for

English Wikipedia. In Proceedings of the Seventh In-

ternational Conference on Language Resources and

Evaluation (LREC'10), Valletta, Malta. European

Maayan Geffet and Ido Dagan. 2005. The distribu-

tional inclusion hypotheses and lexical entailment.

In Proceedings of the 43rd Annual Meeting of the

Association for Computational Linguistics (ACL'05),

pages 107-114, Ann Arbor, Michigan. Association

Irina Higgins, Loic Matthey, Arka Pal, Christopher

Burgess, Xavier Glorot, Matthew Botvinick, Shakir

Mohamed, and Alexander Lerchner. 2017. beta-VAE:

Learning basic visual concepts with a constrained

variational framework. In Proceedings of the 5th In-

ternational Conference on Learning Representations

Lili Kotlerman, Ido Dagan, Idan Szpektor, and Maayan

Zhitomirsky-Geffet. 2010. Directional distributional

similarity for lexical inference. Natural Language

Alessandro Lenci and Giulia Benotto. 2012. Identi-

fying hypernyms in distributional semantic spaces.

In \*SEM 2012: The First Joint Conference on Lexi-

cal and Computational Semantics - Volume 1: Pro-

ceedings of the main conference and the shared task,

and Volume 2: Proceedings of the Sixth Interna-

tional Workshop on Semantic Evaluation (SemEval 2012), pages 75–79, Montréal, Canada. Association

Omer Levy and Yoav Goldberg. 2014. Dependency-

based word embeddings. In Proceedings of the 52nd

Annual Meeting of the Association for Computational

Linguistics (Volume 2: Short Papers), pages 302–308,

Baltimore, Maryland. Association for Computational

Chun Hei Lo, Hong Cheng, Wai Lam, and Guy Emerson.

2023. Functional distributional semantics at scale. In

Proceedings of the 12th Joint Conference on Lexical

and Computational Semantics (\*SEM 2023), pages

423-436, Toronto, Canada. Association for Compu-

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey

Dean. 2013. Efficient estimation of word represen-

tations in vector space. In Proceedings of the 1st

International Conference on Learning Representa-

George A Miller. 1995. WordNet: a lexical database

Ludovica Pannitto, Lavinia Salicchi, and Alessandro

Lenci. 2018. Refining the distributional inclusion

hypothesis for unsupervised hypernym identification.

for English. Communications of the ACM, 38(11):39-

Language Resources Association (ELRA).

for Computational Linguistics.

Engineering, 16(4):359–389.

for Computational Linguistics.

Linguistics.

tational Linguistics.

tions (ICLR), Workshop Track.

(ICLR).

706

- 711 712

714 715

719 720

721

723

727

731

733

737

738

739

740 741

743

744

745

746

747

750

IJCoL. Italian Journal of Computational Linguistics, 4(4-2):45-55. 751

41.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. Pytorch: An imperative style, high-performance deep learning library. In Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc.

752

753

754

755

756

759

760

761

762

763

765

767

768

769

770

771

772

773

775

776

777

778

779

780

781

782

784

785

786

787

789

790

791

793

794

795

796

797

798

799

800

801

802

803

805

806

807

- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 1532–1543.
- Laura Rimell. 2014. Distributional lexical entailment by topic coherence. In Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, pages 511-519, Gothenburg, Sweden. Association for Computational Linguistics.
- Stephen Roller, Katrin Erk, and Gemma Boleda. 2014. Inclusive yet selective: Supervised distributional hypernymy detection. In Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers, pages 1025–1036, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Stephen Roller, Douwe Kiela, and Maximilian Nickel. 2018. Hearst patterns revisited: Automatic hypernym detection from large text corpora. In *Proceedings* of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), pages 358-363, Melbourne, Australia. Association for Computational Linguistics.
- Enrico Santus, Alessandro Lenci, Qin Lu, and Sabine Schulte im Walde. 2014. Chasing hypernyms in vector spaces with entropy. In Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, volume 2: Short Papers, pages 38-42, Gothenburg, Sweden. Association for Computational Linguistics.
- Enrico Santus, Frances Yung, Alessandro Lenci, and Chu-Ren Huang. 2015. EVALution 1.0: an evolving semantic dataset for training and evaluation of distributional semantic models. In Proceedings of the 4th Workshop on Linked Data in Linguistics: Resources and Applications, pages 64-69, Beijing, China. Association for Computational Linguistics.
- Lars Jørgen Solberg. 2012. A corpus builder for Wikipedia. Master's thesis, University of Oslo.
- Kristina Toutanova, Christopher D. Manning, Dan Flickinger, and Stephan Oepen. 2005. Stochastic HPSG parse disambiguation using the Redwoods corpus. Research on Language and Computation, 3(1):83-105.

- 80
- 810 811

813

814 815

816

817

818

819

820

821

822

823

825

826

827

828

829

834

836

837

843

844

Luke Vilnis and Andrew McCallum. 2015. Word representations via Gaussian embedding. In *Proceedings* of the 3rd International Conference on Learning Representations (ICLR).

- Julie Weeds, Daoud Clarke, Jeremy Reffin, David Weir, and Bill Keller. 2014. Learning to distinguish hypernyms and co-hyponyms. In *Proceedings of COLING* 2014, the 25th International Conference on Computational Linguistics: Technical Papers, pages 2249– 2259, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Julie Weeds, David Weir, and Diana McCarthy. 2004. Characterising measures of lexical distributional similarity. In COLING 2004: Proceedings of the 20th International Conference on Computational Linguistics, pages 1015–1021, Geneva, Switzerland. COL-ING.
  - Gisle Ytrestøl, Dan Flickinger, and Stephan Oepen. 2009. Extracting and annotating Wikipedia subdomains: Towards a new eScience community resource. In *Proceedings of the 7th International Workshop on Treebanks and Linguistic Theories*, pages 185–197.
  - Lotfi A. Zadeh. 1965. Fuzzy sets. Information and Control, 8(3):338–353.

## A Derivation of Hypernymy Conditions

Consider (6).  $\forall z$ :

$$t^{(r_H,0)}(z) > t^{(r_h,0)}(z)$$

$$S\left(v^{(r_{H},0)} z + b^{(r_{H},0)}\right) > S\left(v^{(r_{h},0)} z + b^{(r_{h},0)}\right)$$

S is monotonic, so  $\forall z$ :

$$v^{(r_H,0)} z + b^{(r_H,0)} > v^{(r_h,0)} z + b^{(r_h,0)}$$
$$b^{(r_H,0)} - b^{(r_h,0)} > \left(v^{(r_h,0)} - v^{(r_H,0)}\right)^\top z$$

For z on a unit hypercube, this is equivalent to

$$b^{(r_H,0)} - b^{(r_h,0)} > \max_{\|z\|_{\infty}=1} \left( v^{(r_h,0)} - v^{(r_H,0)} \right)^{\top} z$$

Note that  $\arg \max_{z_i} \left( v^{(r_h,0)} - v^{(r_H,0)} \right)^\top z =$  $\operatorname{sgn}(v_i^{(r_h,0)} - v_i^{(r_H,0)})$  where sgn is the sign function. Hence, we have

$$b^{(r_H,0)} - b^{(r_h,0)} > \left\| v^{(r_H,0)} - v^{(r_h,0)} \right\|_{1}$$

For z on a unit hypersphere, this is equivalent to

$$b^{(r_H,0)} - b^{(r_h,0)} > \max_{\|z\|_2 = 1} \left( v^{(r_h,0)} - v^{(r_H,0)} \right)^\top z$$

8 Note that  $\arg \max_{z} \left( v^{(r_{h},0)} - v^{(r_{H},0)} \right)^{\top} z = \frac{v^{(r_{h},0)} - v^{(r_{H},0)}}{\left\| v^{(r_{h},0)} - v^{(r_{H},0)} \right\|_{2}}$ . Hence, we have

$$b^{(r_H,0)} - b^{(r_h,0)} > \left\| v^{(r_H,0)} - v^{(r_h,0)} \right\|_2$$

### **B** Derivation of Transitivity

$$s(r_1, r_2) + s(r_2, r_3)$$
 852

$$= b^{(r_2,0)} - b^{(r_1,0)} - \left\| v^{(r_2,0)} - v^{(r_1,0)} \right\|_p$$
85

$$+ b^{(r_3,0)} - b^{(r_2,0)} - \left\| v^{(r_3,0)} - v^{(r_2,0)} \right\|_p$$
854

$$=b^{(r_3,0)} - b^{(r_1,0)} -$$
85

$$\left(\left\|v^{(r_2,0)} - v^{(r_1,0)}\right\|_p + \left\|v^{(r_3,0)} - v^{(r_2,0)}\right\|_p\right)$$

By the Minkowski inequality, the last term is greater than  $||v^{(r_3,0)} - v^{(r_1,0)}||_p$ . Besides, when  $s(r_1,r_2) > 0$  and  $s(r_2,r_3) > 0$ ,  $s(r_1,r_2) + s(r_2,r_3) > 0$ . Hence,

$$b^{(r_3,0)} - b^{(r_1,0)} - \left\| v^{(r_3,0)} - v^{(r_1,0)} \right\|_p > 0$$
 861

 $s(r_3, r_1) > 0$ 

851

856

857

858

859

860

862

863

864

865

866

867

868

869

870

871

872

873

874

875

876

877

878

879

880

881

882

883

884

885

886

887

888

889

890

891

893

## C Intuition behind Hypernymy Learning by FDS under the DIH

We hypothesize that the way that FDS models are trained allows hypernymy learning under the DIH. During training described in §3.5, the approximate posterior distributions of pixies are first inferred from the observed graph. After variational inference, the semantic functions of the observed predicates are optimized to be true of the inferred pixie distributions. This process is analogous to the following process under a model-theoretic approach: the entities described by a sentence are first identified, and then the truth conditions of predicates over the entities are updated as asserted by the sentence.

Under the DIH, the contexts of nouns are also contexts of their hypernyms. The local predicate– argument information of nouns, i.e. contexts, is thus repeated for their hypernyms for inference during training. Consequently, the semantic functions of hypernyms are trained to return values at least as high as those of their hyponyms over the pixie distributions inferred from the same contexts. The additional contexts appearing exclusively with the hypernyms will further increase the probability of truths of the hypernyms over the pixie space. By (6), hypernymy should thus be learnt under the DIH.

## **D** Training Details

#### **D.1** Hyperparameters and Tuning

For all the experiments, the hyperparameters of the FDS models largely follow that of FDSAS<sub>id</sub> in Lo



Figure 5: Visualization of semantic functions of a run trained on  $H_{\text{chains}}$ . Each plot shows a pixie space in a unit square (unit circle in grey). Each line plots  $t^{(r_i,0)}(z) = 0$  and the arrow points to the pixie subspace where  $t^{(r_i,0)}(z) > 0$ .

et al. (2023) except that we set  $\beta_1$  to 0.5 instead of 0. The consequence is that the inferred pixie distributions during VAE training will be centred closer to the origin. This is motivated by our decision in §3.4 that pixies are only meaningful within the unit hypersphere or hypercube.

Here are the changes exclusive to the experiments on the synthetic data sets. We set K to 1 and perform random negative sampling without weighing by unigram distribution, which trains models maximally using information from the data with minimal assumptions needed for the negative samples. We set the learning rate to 0.01. For experiments on  $H_{\text{chains}}$ , d is set to 2. For  $H_{\text{WN}}$ , d is set to 50. For the remaining topologies, d is set to 10. The models are trained for 2 epochs for  $H_{\text{WN}}$ , and 5000 epochs for the rest.

## D.2 Computational Configurations

912All models are implemented in PyTorch (Paszke913et al., 2019) and trained with distributed data par-914allelism on three NVIDIA GeForce GTX 1080 Ti.915Training a run of FDS or FDS $\forall$  on Wikiwoods takes916about 360 GPU hours.

## **E** Visualization of Semantic Functions

A visualization of results on  $H_{\text{chains}}$  is provided in 918 Fig. 5. As seen in Figs. 5a and 5c, training FDS on 919 the DIH corpus and  $FDS_{\forall}$  on the rDIH corpus both 920 result in four nicely divided pixie subspaces, each 921 for one of the four hypernymy chains, as shown in 922 the plots on the left column. In contrast, applying 923 the other models sometimes gives badly learnt se-924 mantic functions, as shown in Figs. 5b and 5d. For 925 example,  $t^{(r_{12},0)}$  points to the opposite direction of 926  $t^{(r_{10},0)}$  and  $t^{(r_{11},0)}$  in Fig. 5b. 927

917

894