

Modelling Commonsense Commonalities with Multi-Facet Concept Embeddings

Hanane Kteich¹, Na Li³, Usashi Chatterjee², Zied Bouraoui¹, Steven Schockaert²

¹ CRIL CNRS & University of Artois, France ² CardiffNLP, Cardiff University, UK

³ University of Shanghai for Science and Technology, China

{kteich,bouraoui}@cril.fr {chatterjee,schockaerts1}@cardiff.ac.uk

li_na@usst.edu.cn

Abstract

Concept embeddings offer a practical and efficient mechanism for injecting commonsense knowledge into downstream tasks. Their core purpose is often not to predict the commonsense properties of concepts themselves, but rather to identify commonalities, i.e. sets of concepts which share some property of interest. Such commonalities are the basis for inductive generalisation, hence high-quality concept embeddings can make learning easier and more robust. Unfortunately, standard embeddings primarily reflect basic taxonomic categories, making them unsuitable for finding commonalities that refer to more specific aspects (e.g. the colour of objects or the materials they are made of). In this paper, we address this limitation by explicitly modelling the different facets of interest when learning concept embeddings. We show that this leads to embeddings which capture a more diverse range of commonsense properties, and consistently improves results in downstream tasks such as ultra-fine entity typing and ontology completion.

1 Introduction

Many knowledge engineering tasks require knowledge about the meaning of concepts. As a motivating example, let us consider the problem of ontology expansion, which consists in uncovering properties of, and relationships between concepts, given the names of these concepts and an initial knowledge base. Despite the popularity of Large Language Models (LLMs), the use of concept embeddings remains attractive in such settings (Vedula et al., 2018; Li et al., 2019; Malandri et al., 2021; Shi et al., 2023). Indeed, using LLMs directly is often impractical and highly inefficient, as ontologies can involve tens of thousands of concepts. Concept embeddings can also be integrated with structural features more easily (Li et al., 2019), for instance by using them to initialise the node embeddings

of a Graph Neural Network (GNN). Concept embeddings similarly play an important role in many multi-label classification tasks, especially in the zero-shot or few-shot regime (Xing et al., 2019; Yan et al., 2021; Luo et al., 2021; Huang et al., 2022; Ma et al., 2022). As a representative example of such a task, we will consider the problem of ultra-fine entity typing (Choi et al., 2018), which consists in assigning semantic types to mentions of entities, where a total number of around 10K candidate types are considered. In such tasks, the role of pre-trained concept embeddings is to inject prior knowledge about the meaning of the type labels (Xiong et al., 2019; Li et al., 2023a). Note that we cannot straightforwardly accomplish this with LLMs, as they have been found to struggle with information extraction tasks (Han et al., 2023). Moreover, scalability is often an important concern for information extraction systems, which further complicates the use of LLMs.

We take the view that concept embeddings, in the aforementioned applications, are primarily needed to capture commonalities among the concepts involved. For ontology expansion, this is true by definition, since the task explicitly involves identifying sets of concepts that have some property in common. For ultra-fine entity typing, Li et al. (2023a) reported that directly using pre-trained label embeddings was challenging. Instead, they proposed to cluster the set of labels based on pre-trained concept embeddings, and to use the resulting clusters to structure the label space. The idea of using embeddings to structure the label space also lies at the heart of many traditional approaches for zero-shot and few-shot classification.

A key limitation of traditional concept embeddings comes from the fact that they primarily reflect basic taxonomic categories. For instance, the embedding of *banana* is typically similar to that of other fruits, but dissimilar from the embeddings of other yellow things. Some authors have pro-

posed to learn multi-facet embeddings as a way of alleviating these concerns (Rothe and Schütze, 2016; Jain et al., 2018; Alshaiikh et al., 2019, 2020). Essentially, rather than learning a single vector representation of each concept (or entity), they learn a fixed number of different vectors, each focusing on a different facet. However, learning such representations is challenging for two main reasons. First, learning multi-facet representations requires some kind of supervision signal about the facets of interest (Locatello et al., 2019), which is not readily available for many domains. Second, existing approaches consider a fixed set of facets, which makes them unsuitable for open-domain settings. Indeed, the facets of interest strongly depend on the nature of the concepts involved. When modelling food, we may be interested in embeddings that capture their nutritional content. When modelling household appliances, we may want a representation that captures where in the house they are typically found. Rather than using the same set of facets for all concepts, we thus need a more dynamic representation framework.

In this paper, we propose a novel method for learning multi-facet concept embeddings based on two key ideas. First, we rely on ChatGPT¹ to collect a diverse set of (property, facet) pairs, such as (*yellow, colour*), (*found in the kitchen, location*) or (*sweet, taste*), allowing us to treat the problem of learning multi-facet embeddings as a supervised learning problem. Second, rather than learning several independent vector representations, we only learn a single embedding for each concept, treating facets instead as masks on the set of coordinates. This approach offers several modelling advantages, including the fact that facets can have a hierarchical structure (e.g. *colour* is a sub-facet of *appearance*) and the fact that we do not have to tune the number and dimension of the facets a priori. Specifically, we train three BERT (Devlin et al., 2019) encoders: one encoder to map concepts onto their embedding, one encoder to map properties onto their embedding, and one encoder to map properties onto the embedding of the corresponding facet. We show that these encoders can be effectively trained using only training data obtained from ChatGPT, although the best results are obtained by augmenting this training data with examples from ConceptNet².

¹<https://openai.com>

²<https://conceptnet.io>

2 Related Work

Concept Embeddings The idea that language models of the BERT family can be used for learning concept embeddings has been studied extensively. Some approaches simply use the name of the concept as input to the BERT encoder, possibly together with a short prompt (Bommasani et al., 2020; Vulić et al., 2021; Liu et al., 2021a; Gajbhiye et al., 2022). Other approaches instead use contextualised representations from sentences mentioning the concept, selected from some corpus (Ethayarajh, 2019; Bommasani et al., 2020; Vulić et al., 2020; Liu et al., 2021b; Li et al., 2023b). These approaches have been developed with different motivations in mind. One common motivation is to learn something about the language model itself by inspecting the resulting concept embeddings, such as biases (Bommasani et al., 2020) or the model’s grasp of lexical semantics (Ethayarajh, 2019; Vulić et al., 2020). Other authors have rather focused on the use of embeddings for predicting semantic properties of concepts (Gajbhiye et al., 2022; Li et al., 2023b; Rosenfeld and Erk, 2023). Our paper can be seen as a continuation of this latter research line, where we aim to improve the range of properties that can be captured by concept embeddings through the use of facet embeddings.

Commonalities Gajbhiye et al. (2023) recently argued that the main purpose of concept embeddings, when it comes to downstream applications, is usually to identify what different concepts have in common. Specifically, given a set of concepts, they first used the corresponding concept embeddings to predict a set of properties for each concept. The resulting predictions were then filtered using a Natural Language Inference (NLI) model. Finally, properties that were found for at least two concepts were identified as shared properties. They showed, on the task of ultra-fine entity typing, that by augmenting the training data with these shared properties, models were able to generalise better. This idea also relates to the notion of conceptualisation (He et al., 2022; Wang et al., 2023b). Essentially, the latter works have suggested to augment commonsense knowledge graphs by generalising the concepts involved. This often involves replacing a specific concept (e.g. a football game) by a description referring to some salient property (e.g. a relaxing event). Wang et al. (2023a) showed that the resulting generalisations of commonsense knowledge graphs were useful for zero-shot com-

nonsense question answering. The aforementioned methods all rely on the availability of a set of properties (or hypernyms) that can be used to generalise a given set of concepts. In practice, however, it is hard to obtain comprehensive property sets, which means that many commonalities may not be discovered. Moreover, certain commonalities are hard to describe, even though they intuitively make sense.³ To avoid such limitations, we identify commonalities by clustering concept embeddings.

Multi-Facet Embeddings The idea of capturing different facets of meaning has been studied in the context of disentangled representation learning, especially in computer vision (Chen et al., 2016; Higgins et al., 2017; Kim and Mnih, 2018; Chen et al., 2018). When it comes to learning disentangled representation of text, He et al. (2017) proposed a method for learning aspect embeddings in the context of sentiment analysis, whereas several authors have proposed multi-facet document embeddings (Jain et al., 2018; Risch et al., 2021; Kohlmeyer et al., 2021). Rothe and Schütze (2016) suggested that word embeddings could be decomposed into meaningful subspaces, which essentially correspond to facets. Most similar to our work, Alshaikh et al. (2019) proposed a method for decomposing a domain-specific concept embedding space into subspaces capturing different facets. To find this decomposition, they relied on the assumption that properties belonging to the same facet tend to have similar word embeddings. Finally, Alshaikh et al. (2020) proposed a mixture-of-experts model to learn multi-facet concept embeddings directly, using a variant of GloVe (Pennington et al., 2014).

3 Proposed Approach

We propose a bi-encoder based concept embedding model which is capable of representing concepts w.r.t. different facets. The key stumbling block in previous work on learning multi-facet embeddings has been the difficulty in acquiring a meaningful supervision signal about which properties belong to the same facet (e.g. that *large* and *small* both refer to *size*). As explained in Section 3.1, we can now overcome this difficulty by collecting property-facet pairs from LLMs. Our proposed model itself is described in Section 3.2. Finally, we explain how facet-specific embeddings can be extracted once

³As a simple toy example, among the set {*cat*, *dog*, *goldfish*, *rabbit*} the concepts *cat* and *dog* stand out as similar, even though they are not the only pets nor the only mammals.

the model has been trained (Section 3.3).

3.1 Obtaining Training Data

We need two types of training examples for our model: concept-property judgements (e.g. *banana* has the property *rich in potassium*) and property-facet judgments (e.g. *rich in potassium* refers to *nutritional content*). We use two sources for obtaining these examples: ChatGPT and ConceptNet.

ChatGPT We use the dataset of 109K concept-property judgments that were obtained from ChatGPT by Chatterjee et al. (2023).⁴ To obtain property-facet pairs, we proceed in a similar way, although obtaining suitable information about facets turned out to be more challenging. We obtained the best results with the following prompt, which does not ask about facets explicitly. Instead, we ask about concept-property pairs, but use a format which requires the model to specify the facet of each property that is generated:

*I am interested in knowing common properties that are satisfied by different concepts. 1. **Sound: loud - thunder, jet engine, siren** 2. **Temperature: cold - ice, refrigerator, Antarctica** 3. **Colour: orange - mandarin, basketball, clownfish** 4. **Shape: round - sun, orange, ball** 5. **Purpose: used for cleaning - broom, lemon, soap** 6. **Location: located in the ocean - sand, whale, corals.** Please provide me with a list of 30 such examples.*

We used this request with the same prompt 10 times. After this, we change the examples that are given (shown in bold above) and repeat. We manually processed the responses to standardise facet spellings and removed duplicates. For instance, facets were sometimes generated in plural (e.g. *colors* rather than *color*), or the same facet was generated with different spellings (e.g. *color* and *colour*). Even when changing the examples in the prompt after every 10 requests, the number of unique facet-property pairs that were generated saturated relatively quickly. In total, we obtained 828 unique facet-property pairs, covering 127 unique facets.

ConceptNet Starting from a ConceptNet 5 dump⁵, we first selected the English language

⁴Dataset available from https://github.com/ExperimentsLLM/EMNLP2023_PotentialOfLLM_LearningConceptualSpace.

⁵<https://github.com/commonsense/conceptnet5/wiki/Downloads>

triples. Given a triple such as *(boat, at location, sea)* we create a corresponding concept-property pair *(boat, at location sea)* and a property-facet pair *(at location sea, at location)*. In other words, the ConceptNet relations are treated as facets, and properties are obtained by combining a relation with a tail concept. Not all ConceptNet relations are suitable for this purpose. We specifically used: RelatedTo, FormOf, IsA, UsedFor, AtLocation, CapableOf, HasProperty, HasA, InstanceOf and MadeOf. Furthermore, when creating properties, we only consider tail concepts that appear at least 10 times. We thus end up with 884 distinct properties, 10 facets, 18505 concepts, 884 property-facet pairs and 36955 concept-property pairs.

3.2 Model Formulation

Let us write \mathcal{D}_{cp} for the set of (concept, property) pairs that are available for training. Similarly, we write \mathcal{D}_{pf} for the set of available (property, facet) pairs. We build on the following bi-encoder loss from (Gajbhiye et al., 2022):

$$\begin{aligned} \mathcal{L} = & - \sum_{(c,p) \in \mathcal{D}_{cp}} \log \sigma(\text{Con}(c) \cdot \text{Prop}(p)) \\ & - \sum_{(c,p) \in \mathcal{N}_{cp}} \log(1 - \sigma(\text{Con}(c) \cdot \text{Prop}(p))) \end{aligned}$$

where \mathcal{N}_{cp} is a set of negative examples. Specifically, for each positive example (c, p) , five negative examples (c, p') are obtained by replacing p by another property p' . The concept embedding $\text{Con}(c)$ and property embedding $\text{Prop}(p)$ are obtained by two separate BERT encoders. The concept encoder Con uses a prompt of the form $\langle \text{Concept} \rangle$ means [MASK]. The property encoder Prop uses the same prompt. In both cases, the embeddings are obtained from the final-layer embedding of the [MASK] token. However, the concept embedding $\text{Con}(c)$ is normalised (w.r.t. the Euclidean norm) whereas the property embedding $\text{Prop}(p)$ is not.

Previous work on multi-facet embeddings has relied on learning multiple concept embeddings, where each concept embedding only captures a subset of all properties (Alshaiikh et al., 2020). This approach has a number of drawbacks, however. For instance, it relies on the idea that each facet is represented using the same number of dimensions, implicitly assumes that the overall number of facets is relatively small, and that the facets are independent from each other. This is particularly problematic in open-domain settings, where a wide range of

facets may need to be considered, certain facets only make sense for some concepts (e.g. nutritional value only makes sense for food) and facets often have a hierarchical structure (e.g. colour is a sub-facet of appearance). Therefore, instead of learning multiple concept embeddings, we instead interpret facets as masks on concept embeddings.

Specifically, we train a third BERT encoder, Facet , which also takes the property p as input and again uses the same prompt. The idea is that $\text{Facet}(p)$ indicates which coordinates of the concept embeddings are most relevant when modelling the property p . We define the masked embedding of concept c w.r.t. some property p as follows:

$$\text{MC}(c, p) = \frac{\text{Con}(c) \odot \text{Facet}(p)}{\|\text{Con}(c) \odot \text{Facet}(p)\|}$$

where we write \odot for the component-wise product. We essentially keep the same bi-encoder model but instead rely on these masked concept embeddings:

$$\begin{aligned} \mathcal{L}_1 = & - \sum_{(c,p) \in \mathcal{D}_{cp}} \log \sigma(\text{MC}(c, p) \cdot \text{Prop}(p)) \\ & - \sum_{(c,p) \in \mathcal{N}_{cp}} \log(1 - \sigma(\text{MC}(c, p) \cdot \text{Prop}(p))) \end{aligned}$$

Without further supervision, the facet encoder does not learn meaningful facets. Therefore, we use the (property, facet) examples from \mathcal{D}_{pf} to ensure that properties which belong to the same facet have a similar facet embedding. For a given facet f , we write \mathcal{P}_f for the set of properties that we know to belong to this facet, i.e. $\mathcal{P}_f = \{p \mid (p, f) \in \mathcal{D}_{pf}\}$. We use the InfoNCE loss:

$$\mathcal{L}_2 = - \sum_f \sum_{p, q \in \mathcal{P}_f} \log \frac{\exp\left(\frac{\cos(F(p), F(q))}{\tau}\right)}{\sum_r \exp\left(\frac{\cos(F(p), F(r))}{\tau}\right)}$$

where we abbreviated $\text{Facet}(p)$ as $F(p)$, the sum in the denominator ranges over $r \in \{q\} \cup \{p \mid (p, f) \notin \mathcal{D}_{pf}\}$, and the temperature $\tau > 0$ is a hyperparameter. The InfoNCE loss encourages properties which belong to the same facet to have facet embeddings that are more similar to each other than to the facet embeddings of properties which do not. The overall model is trained by optimising the loss $\mathcal{L}_1 + \mathcal{L}_2$.

3.3 Extracting Facet-Specific Representations

The model from Section 3.2 can be used in several ways. First, we can simply use the concept embeddings $\text{Con}(c)$ to represent each concept c .

In this case, the purpose of having facets is to ensure that the concept embeddings capture a broader range of properties, but we only consider these facets during training. We will refer to this approach as **ConEmb-F**. The concept embeddings from the standard bi-encoder, without facet embeddings, will be referred to as **ConEmb**.

In some applications, concept embeddings are used for clustering concepts. The purpose of multi-facet embeddings is to ensure that different kinds of clusters can be found. In such settings, we extract different facet-specific concept embeddings from the model. Specifically, let $\mathcal{P} = \{p_1, \dots, p_m\}$ be the set of properties of interest. For each property p_i we have a corresponding facet vector $\mathbf{f}_i = \text{Facet}(p_i)$. We use the K-means algorithm to cluster these facet vectors into clusters $\mathcal{X}_1, \dots, \mathcal{X}_k$ and treat each of these clusters as facet. We associate each concept c with k facet-specific representations $\mathbf{c}_1, \dots, \mathbf{c}_k$, defined as:

$$\mathbf{c}_j = \frac{\text{Con}(c) \odot \left(\sum_{p_i \in \mathcal{X}_j} \mathbf{f}_i \right)}{\left\| \text{Con}(c) \odot \left(\sum_{p_i \in \mathcal{X}_j} \mathbf{f}_i \right) \right\|} \quad (1)$$

The representations obtained by this approach depend on how the set of properties \mathcal{P} is chosen. For our experiments, we simply set \mathcal{P} to be the set of all properties that appear in our training set.

4 Experiments

We analyse the effectiveness of the proposed multi-facet concept embedding model. We intrinsically evaluate the embeddings on predicting commonsense properties (Section 4.1) and outlier detection (Section 4.2). We furthermore consider two downstream applications: ontology completion (Section 4.3) and ultra-fine entity typing (Section 4.4).⁶

4.1 Predicting Commonsense Properties

The use of facets should lead to concept embeddings that capture a wider range of properties. To test this hypothesis, we consider the task of commonsense property prediction, which we treat as a binary classification problem: given a concept and a commonsense property, decide whether the property is satisfied by the concept or not. The difficulty of this task depends on how the training-test split is constructed. One strategy, called **concept split**, ensures that the concepts appearing in the

training and test data are different, but the properties are not. Gajbhiye et al. (2022) found that simple nearest neighbour strategies can do well on this task, meaning that this variant does not adequately assess whether the concept embeddings capture commonsense knowledge. For this reason, they proposed a **property split**, where the properties appearing in training and test are different, but the concepts are the same. Finally, they also considered a **C+P** (concept+property) split, where both the concepts and properties are different in training and test. In all cases, we first pre-train the encoders Con, Prop and Facet on the ChatGPT and/or ConceptNet training data (see Section 3.1), before fine-tuning on the training splits of the property prediction benchmarks.⁷

Table 1 summarises the results for all three settings, using two standard benchmarks for commonsense property prediction: the extension of the McRae property norms dataset (McRae et al., 2005) that was introduced by Forbes et al. (2019) and the augmented version of CSLB⁸ introduced by Misra et al. (2022). Our main baseline is the bi-encoder model from Gajbhiye et al. (2022), shown as **BiEnc**, which also forms the basis for our facet-based model. We report the results from Gajbhiye et al. (2022), which are for a model that was pre-trained on Microsoft Concept Graph (Ji et al., 2019) and GenericsKB (Bhakthavatsalam et al., 2020), as well as results for models that we trained on the ChatGPT and ConceptNet training sets (see Section 3.1). Finally, we show the result of our full model (i.e. with the facet encoder), shown as **BiEnc-F**. We compare four different encoders: BERT-base, BERT-large, RoBERTa-base and RoBERTa-large.

The results clearly show the benefit of using facets, as the BiEnc-F models consistently and substantially outperform the BiEnc baselines. Comparing the different training sets, the ChatGPT training examples are more effective than the ConceptNet examples, but the best results are obtained when both sources are combined. Among the different LMs, BERT-large achieves the best results.

Based on the results from this experiment, for the remaining experiments, we will focus on the model based on BERT-large, which is trained on both the ChatGPT and ConceptNet training examples.

⁶Our datasets, pre-trained models and implementation are available at https://github.com/hananeekh/facets_concepts_embeddings

⁷Detailed training details can be found in the appendix.

⁸<https://cs1b.psychol.cam.ac.uk/propnorms>

	LM	Train properties (\mathcal{D}_{cp})	Train facets (\mathcal{D}_{pf})	McRae			CSLB		
				Con	Prop	C+P	Con	Prop	C+P
BiEnc*	BB	MSCG	-	79.8	49.6	44.5	54.5	39.1	32.6
BiEnc*	BL	MSCG	-	80.5	49.3	45.5	57.7	41.8	36.4
BiEnc*	RB	MSCG	-	75.6	42.4	38.1	49.9	36.4	24.3
BiEnc*	RL	MSCG	-	80.1	46.5	42.5	59.0	42.5	36.0
BiEnc	BL	CN	-	78.0	56.7	51.8	61.4	49.6	50.0
BiEnc	BL	ChatGPT	-	80.5	57.3	56.6	65.1	56.5	52.7
BiEnc	BL	ChatGPT+CN	-	81.7	62.1	59.5	67.8	59.6	53.1
BiEnc	BB	ChatGPT+CN	-	76.2	60.6	58.4	66.9	56.6	51.8
BiEnc	RB	ChatGPT+CN	-	75.8	60.1	58.2	66.1	56.3	51.7
BiEnc	RL	ChatGPT+CN	-	80.8	61.7	59.3	67.2	58.8	52.7
BiEnc-F	BL	ChatGPT+CN	CN	84.3	63.5	57.7	69.4	61.0	59.9
BiEnc-F	BL	ChatGPT+CN	ChatGPT	84.3	64.9	65.5	69.5	61.6	61.9
BiEnc-F	BL	ChatGPT+CN	ChatGPT+CN	86.2	65.9	67.0	70.3	63.6	63.0
BiEnc-F	BB	ChatGPT+CN	ChatGPT+CN	82.1	63.0	61.2	65.3	60.2	59.9
BiEnc-F	RB	ChatGPT+CN	ChatGPT+CN	81.5	62.3	60.8	65.0	59.6	61.3
BiEnc-F	RL	ChatGPT+CN	ChatGPT+CN	85.6	65.1	65.9	69.2	63.1	62.8

Table 1: Results for commonsense property prediction in terms of F1 (%). Results marked with * were taken from Gajbhiye et al. (2022). MSCG corresponds to the training set from Gajbhiye et al. (2022); ChatGPT and CN (ConceptNet) refer to the training sets that were described in Section 3.1. We evaluate using BERT-base-uncased (BB), BERT-large-uncased (BL), RoBERTa-base (RB) and RoBERTa-large (RL).

	ConEmb	ConEmb-F	MultiConEmb
Dangerous	9	3	13
Edible	23	26	67
Flies	23	34	77
Hot	7	11	40
Lives in water	33	48	83
Produces noise	13	13	48
Sharp	33	30	67
Used by children	8	9	17
Used for cooking	59	53	93
Worn on feet	18	13	54

Table 2: Results for outlier detection (exact match %).

4.2 Outlier Detection

To evaluate whether our facet-based concept embeddings can help us to identify commonalities, we consider the task of outlier detection (Camacho-Collados and Navigli, 2016; Blair et al., 2017; Brink Andersen et al., 2020). In each instance of this task, we are given a set of concepts (or entities). Some of these concepts have a particular property in common, and the task consists in identifying these concepts (without being given any information about the shared property itself). This task has traditionally been used as an intrinsic benchmark for evaluating word embeddings.

Dataset Existing benchmarks mostly focus on broad taxonomic categories, whereas we are specifically interested in identifying shared commonsense properties. We therefore constructed a new out-

lier detection benchmark based on the extended McRae dataset from Forbes et al. (2019). To create an outlier detection problem instance, we first select a property from this dataset (e.g. *dangerous*) as well as 3 concepts which have this property and 7 outlier concepts which do not. Many of the properties in the McRae dataset are not suitable for our benchmark, either because they correspond to taxonomic categories (e.g. *an animal* or *edible*) or because they are too ambiguous or noisy (e.g. *accordion*, *car* and *escalator* are described as having the property *fun*, but *airplane* is not). Therefore, we manually selected 10 properties which do not suffer from these limitations. For each property, we manually clustered the concepts with this property into broad taxonomic groups.⁹ When selecting the 3 positive concepts, for a given instance, we ensure that all three examples come from a different group. When selecting the 7 outliers, we check that they do not share any properties. Specifically, we ensure that any two of the outliers do not have any properties in common in the extended McRae dataset, in ConceptNet or in Ascent++¹⁰. For each property, we sample 100 problem instances, following this process. We report the results in terms of exact match, i.e. the prediction for a given instance is labelled as correct if the three positive examples were correctly identified. We report the percentage

⁹The resulting clusters can be found in the appendix.

¹⁰<https://ascentpp.mpi-inf.mpg.de>

	Wine	Econ	Olym	Tran	SUMO
GloVe*	14.2	14.1	9.9	8.3	34.9
Skipgram*	13.8	13.5	8.3	7.2	33.4
Numberbatch*	25.6	26.2	26.8	16.0	47.3
MirrorBERT*	22.5	23.8	20.9	12.7	40.1
MirrorWiC*	24.7	24.9	22.1	13.9	46.9
ConCN*	31.3	32.4	29.7	20.9	52.6
ConEmb	30.8	30.5	28.6	19.8	51.3
ConEmb-F	31.2	31.8	30.4	20.9	51.7
Clu (ConCN)	35.3	33.1	32.5	21.6	52.2
Clu (ConEmb-F)	36.9	34.2	34.6	22.1	53.3
MClu (ConCN)	39.8	35.9	32.6	22.7	54.2
MClu (ConEmb-F)	39.9	36.3	32.9	23.1	55.4

Table 3: Results for ontology completion in terms of F1 (%). Results marked with * were taken from Li et al. (2023b). ConEmb and ConEmb-F were trained using both the ChatGPT and ConceptNet training sets.

of correctly labelled instances, for each property.

Methods We compare three strategies for detecting outliers. For the method denoted **ConEmb**, we use the ConEmb embeddings as follows. For each concept c , we find the second and third nearest neighbour. Let us denote these as c_2 and c_3 . If c is a positive example, $\cos(\text{Con}(c), \text{Con}(c_2))$ should be high and $\cos(\text{Con}(c), \text{Con}(c_3))$ should be low. We thus score each concept as $\text{score}(c) = \cos(\text{Con}(c), \text{Con}(c_2)) - \cos(\text{Con}(c), \text{Con}(c_3))$. As positive examples, we then select the concept with the highest score along with its two nearest neighbours. The method denoted **ConEmb-F** uses the same strategy, but instead uses the ConEmb-F embeddings. Finally, when using the method denoted **MultiConEmb**, we first obtain 10 facet-specific embeddings of each concept, using (1). We then first apply the same method as before to each of the 10 facet-specific embedding spaces. Finally, we select the prediction for the facet where the score of the highest-scoring concept was maximal.

Results The results are summarised in Table 2. MultiConEmb, which exploits facet-specific representations, substantially outperforms the baselines. The performance of ConEmb and ConEmb-F is comparable, which is as expected: even though ConEmb-F was trained using facets, this method represents concepts as single vectors, and the similarities between these concept vectors still mostly reflect taxonomic relatedness.

4.3 Ontology Completion

Ontologies use rules to encode how the concepts of a given domain are related. They generalise tax-

onomies by allowing the use of logical connectives to encode these relationships. Li et al. (2019) introduced a framework for predicting missing rules in ontologies using a Graph Neural Network. The nodes of the considered graph correspond to the concepts from the ontology, and the input representations are pre-trained concept embeddings. Recent work has used this model to evaluate concept embeddings, as its overall performance is sensitive to the quality of the input representations (Li et al., 2023b). The intuition underpinning the model is closely aligned with the idea of modelling concept commonalities. Essentially, if the ontology contains the rules¹¹ $X_1 \sqsubseteq Y, \dots, X_k \sqsubseteq Y$ and we know from the pre-trained embeddings that X_{k+1} is similar to X_1, \dots, X_k then it is plausible that the rule $X_{k+1} \sqsubseteq Y$ is valid as well.

We test the effectiveness of our model in two ways. First, we use the **ConEmb-F** concept embeddings as input features, which allows for a direct comparison with the effectiveness of other concept embedding models. In this case, the use of facets only affects how the concept embeddings are learned. As a second strategy, referred to as **MClu**, we first obtain 10 facet-specific embeddings of all the concepts, using (1). We then cluster the concepts in each of the facet-specific concept embeddings separately. For this step, we have relied on affinity propagation. This results in 10 different clusterings of the concepts. For each cluster \mathcal{C} in each of these clustering we add a fresh concept $Y_{\mathcal{C}}$ to the ontology, and for every concept X in \mathcal{C} , we add the rule $X \sqsubseteq Y_{\mathcal{C}}$. We then apply the standard GNN model from Li et al. (2019) to the resulting extended ontology. As a baseline, we also apply the clustering strategy to the ConEmb concept embeddings, which we refer to as **Clu**. Note that in this case there is only one clustering. Note that when **MClu** or **Clu** are used, we still need to use concept embeddings to use as input features. We show results with the ConCN embeddings from Li et al. (2023b) and for our ConEmb-F embeddings.

The results in Table 3 show that the ConEmb-F input embeddings consistently outperform the ConEmb vectors. Moreover, they achieve a performance which is similar to that of the ConCN embeddings, which is the current state-of-the-art. Moreover, the proposed clustering strategies, which have not previously been considered for ontology

¹¹In description logic syntax, $X \sqsubseteq Y$ means that every instance of the concept X is also an instance of the concept Y , i.e. it represents the rule “if X then Y ”.

	F1
Base model [†]	49.2
Properties*	50.9
Clu (ConCN) [†]	50.4
Clu (ConEmb)*	50.6
Clu (ConEmb-F)	50.8
Clu (ConCN) + properties*	50.9
Clu (ConEmb) + properties*	51.1
MClu	51.3

Table 4: Results for ultra-fine entity typing, using a BERT-base entity encoder with augmented label sets. Results marked with * were taken from Gajbhiye et al. (2023); results marked with [†] were taken from Li et al. (2023a).

completion, are highly effective. The MClu variant outperforms Clu in all but one case, which shows the benefit of explicitly considering facet-specific embeddings. We can also see that the ConEmb-F embeddings as input features perform better than the ConCN embeddings, when used in combination with the clustering strategies.

4.4 Ultra-Fine Entity Typing

We consider the task of ultra-fine entity typing (Choi et al., 2018), which was also used by Gajbhiye et al. (2023) to demonstrate the usefulness of modelling concept commonalities. Given a sentence in which an entity mention is highlighted, the task consists in assigning labels that describe the semantic type of the entity. The task is formulated as a multi-label classification problem with around 10K candidate labels. Many of the candidate labels only have a small number of occurrences in the training data. This makes it paramount to rely on some kind of pre-trained knowledge about the meaning of the labels. Li et al. (2023a) proposed a simple but surprisingly effective strategy: use pre-trained concept embeddings to cluster the labels and augment the training set with labels that refer to these clusters. For instance, if a training example is labelled with label l and this label belongs to cluster c then they add the synthetic label “cluster c ” to this training example. This intuitively teaches the model which labels are semantically related, as the training objective encourages instances which are labelled with “cluster c ” to be linearly separated from other instances. Gajbhiye et al. (2023) improved on this strategy by instead identifying commonsense properties that were satisfied by the different concepts/labels, and by augmenting the training examples with these properties, instead of

synthetic cluster labels. This use of shared properties has the advantage that a broader range of commonalities can be identified, whereas clustering standard concept embeddings leads to clusters that only reflect standard taxonomic categories. Our hypothesis is that we can achieve the same benefits by clustering our multi-facet representations, and that the use of clusters can potentially lead us to capture finer-grained commonalities.

Table 4 summarises the results. All results were obtained using the DenoiseFET model from Pan et al. (2022). The *base model* in Table 4 shows the results if we use this model without augmenting the training labels. *Properties* refers to the strategy from Gajbhiye et al. (2023), which adds labels corresponding to shared properties, while **Clu** is the strategy from Li et al. (2023a), which adds labels corresponding to clusters. We show results for this clustering strategy with three different concept embeddings: the ConCN embeddings from Li et al. (2023b) as well as ConEmb and ConEmb-F. The approach where we use clusterings from different facet-specific embeddings is shown as **MClu**. We can see that MClu achieves the best results, which confirms the usefulness of facet-specific representations for this task. When using the Clu strategy, ConEmb-F also slightly outperforms ConEmb.

5 Conclusions

Many applications rely on background knowledge about the meaning of concepts. What is needed often boils down to knowledge about the commonalities between different concepts, as this forms the basis for inductive generalisation. Clustering pre-trained concept embeddings has been proposed in previous work as a viable strategy for modelling such commonalities. However, the resulting clusters primarily capture taxonomic categories, while commonalities that depend on various commonsense properties are essentially ignored. In this paper, we proposed a simple strategy for obtaining more diverse representations, by taking into account different facets of meaning when training the concept encoder. We found that the resulting concept representations lead to consistently better results, across all the considered tasks.

Acknowledgments This work was supported by EPSRC grants EP/V025961/1 and EP/W003309/1, ANR-22-CE23-0002 ERIANA and HPC resources from GENCI-IDRIS (Grant 2023-[AD011013338R1]).

Limitations

Our approach relies on encoders from the BERT family, which are much smaller than recent language models. We did some initial experiments with Llama 2, but were not successful in obtaining better-performing concept embeddings with this model. While it seems likely that future work will reveal more effective strategies for using larger models, our use of BERT still has the advantage that we can efficiently encode a large number of labels, which remains important for applications such as extreme multi-label text classification.

We have only looked at modelling common-sense properties of concepts. Modelling facets of meaning is intuitively also important for modelling named entities (e.g. for entity linking) and for sentence/document embeddings (e.g. for retrieval). An analysis of facet-based models for such applications is left as a topic for future work.

References

- Rana Alshaikh, Zied Bouraoui, Shelan Jeawak, and Steven Schockaert. 2020. [A mixture-of-experts model for learning multi-facet entity embeddings](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5124–5135, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Rana Alshaikh, Zied Bouraoui, and Steven Schockaert. 2019. [Learning conceptual spaces with disentangled facets](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 131–139, Hong Kong, China. Association for Computational Linguistics.
- Sumithra Bhakthavatsalam, Chloe Anastasiades, and Peter Clark. 2020. [Genericskb: A knowledge base of generic statements](#). *CoRR*, abs/2005.00660.
- Philip Blair, Yuval Merhav, and Joel Barry. 2017. [Automated generation of multilingual clusters for the evaluation of distributed representations](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*. OpenReview.net.
- Rishi Bommasani, Kelly Davis, and Claire Cardie. 2020. [Interpreting Pretrained Contextualized Representations via Reductions to Static Embeddings](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4758–4781, Online. Association for Computational Linguistics.
- Jesper Brink Andersen, Mikkel Bak Bertelsen, Mikkel Hørby Schou, Manuel R. Ciosici, and Ira Assent. 2020. [One of these words is not like the other: a reproduction of outlier identification using non-contextual word representations](#). In *Proceedings of the First Workshop on Evaluation and Comparison of NLP Systems*, pages 120–130, Online. Association for Computational Linguistics.
- José Camacho-Collados and Roberto Navigli. 2016. [Find the word that does not belong: A framework for an intrinsic evaluation of word vector representations](#). In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*, pages 43–50, Berlin, Germany. Association for Computational Linguistics.
- Usashi Chatterjee, Amit Gajbhiye, and Steven Schockaert. 2023. [Cabbage sweeter than cake? analysing the potential of large language models for learning conceptual spaces](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 11836–11842, Singapore. Association for Computational Linguistics.
- Tian Qi Chen, Xuechen Li, Roger B. Grosse, and David Duvenaud. 2018. [Isolating sources of disentanglement in variational autoencoders](#). In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 2615–2625.
- Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. 2016. [Infogan: Interpretable representation learning by information maximizing generative adversarial nets](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 2172–2180.
- Eunsol Choi, Omer Levy, Yejin Choi, and Luke Zettlemoyer. 2018. [Ultra-fine entity typing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 87–96, Melbourne, Australia. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Kawin Ethayarajh. 2019. [How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 55–65, Hong Kong, China. Association for Computational Linguistics.

- Maxwell Forbes, Ari Holtzman, and Yejin Choi. 2019. [Do neural language representations learn physical commonsense?](#) In *Proceedings of the 41th Annual Meeting of the Cognitive Science Society, CogSci 2019: Creativity + Cognition + Computation, Montreal, Canada, July 24-27, 2019*, pages 1753–1759. cognitivesciencesociety.org.
- Amit Gajbhiye, Zied Bouraoui, Na Li, Usashi Chatterjee, Luis Espinosa-Anke, and Steven Schockaert. 2023. [What do deck chairs and sun hats have in common? uncovering shared properties in large concept vocabularies.](#) In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10587–10596, Singapore. Association for Computational Linguistics.
- Amit Gajbhiye, Luis Espinosa-Anke, and Steven Schockaert. 2022. [Modelling commonsense properties using pre-trained bi-encoders.](#) In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3971–3983, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Ridong Han, Tao Peng, Chaohao Yang, Benyou Wang, Lu Liu, and Xiang Wan. 2023. [Is information extraction solved by chatgpt? an analysis of performance, evaluation criteria, robustness and errors.](#) *CoRR*, abs/2305.14450.
- Mutian He, Tianqing Fang, Weiqi Wang, and Yangqiu Song. 2022. [Acquiring and modelling abstract commonsense knowledge via conceptualization.](#) *CoRR*, abs/2206.01532.
- Ruidan He, Wee Sun Lee, Hwee Tou Ng, and Daniel Dahlmeier. 2017. [An unsupervised neural attention model for aspect extraction.](#) In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 388–397, Vancouver, Canada. Association for Computational Linguistics.
- Irina Higgins, Loïc Matthey, Arka Pal, Christopher P. Burgess, Xavier Glorot, Matthew M. Botvinick, Shakir Mohamed, and Alexander Lerchner. 2017. [beta-vae: Learning basic visual concepts with a constrained variational framework.](#) In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- James Y. Huang, Bangzheng Li, Jiashu Xu, and Muhao Chen. 2022. [Unified semantic typing with meaningful label inference.](#) In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2642–2654, Seattle, United States. Association for Computational Linguistics.
- Sarthak Jain, Edward Banner, Jan-Willem van de Meent, Iain J. Marshall, and Byron C. Wallace. 2018. [Learning disentangled representations of texts with application to biomedical abstracts.](#) In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4683–4693, Brussels, Belgium. Association for Computational Linguistics.
- Lei Ji, Yujing Wang, Botian Shi, Dawei Zhang, Zhongyuan Wang, and Jun Yan. 2019. [Microsoft concept graph: Mining semantic concepts for short text understanding.](#) *Data Intell.*, 1(3):238–270.
- Hyunjik Kim and Andriy Mnih. 2018. [Disentangling by factorising.](#) In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 2654–2663. PMLR.
- Lasse Kohlmeyer, Tim Repke, and Ralf Krestel. 2021. [Novel views on novels: Embedding multiple facets of long texts.](#) In *WI-IAT '21: IEEE/WIC/ACM International Conference on Web Intelligence, Melbourne VIC Australia, December 14 - 17, 2021*, pages 670–675. ACM.
- Na Li, Zied Bouraoui, and Steven Schockaert. 2019. [Ontology completion using graph convolutional networks.](#) In *The Semantic Web - ISWC 2019 - 18th International Semantic Web Conference, Auckland, New Zealand, October 26-30, 2019, Proceedings, Part I*, volume 11778 of *Lecture Notes in Computer Science*, pages 435–452. Springer.
- Na Li, Zied Bouraoui, and Steven Schockaert. 2023a. [Ultra-fine entity typing with prior knowledge about labels: A simple clustering based strategy.](#) In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 11744–11756, Singapore. Association for Computational Linguistics.
- Na Li, Hanane Kteich, Zied Bouraoui, and Steven Schockaert. 2023b. [Distilling semantic concept embeddings from contrastively fine-tuned language models.](#) In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2023, Taipei, Taiwan, July 23-27, 2023*, pages 216–226. ACM.
- Fangyu Liu, Ivan Vulić, Anna Korhonen, and Nigel Collier. 2021a. [Fast, effective, and self-supervised: Transforming masked language models into universal lexical and sentence encoders.](#) In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1442–1459, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Qianchu Liu, Fangyu Liu, Nigel Collier, Anna Korhonen, and Ivan Vulić. 2021b. [MirrorWiC: On eliciting word-in-context representations from pretrained language models.](#) In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 562–574, Online. Association for Computational Linguistics.
- Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Rätsch, Sylvain Gelly, Bernhard Schölkopf, and

- Olivier Bachem. 2019. [Challenging common assumptions in the unsupervised learning of disentangled representations](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 4114–4124. PMLR.
- Qiaoyang Luo, Lingqiao Liu, Yuhao Lin, and Wei Zhang. 2021. [Don't miss the labels: Label-semantic augmented meta-learner for few-shot text classification](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2773–2782, Online. Association for Computational Linguistics.
- Jie Ma, Miguel Ballesteros, Srikanth Doss, Rishita Anubhai, Sunil Mallya, Yaser Al-Onaizan, and Dan Roth. 2022. [Label semantics for few shot named entity recognition](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1956–1971, Dublin, Ireland. Association for Computational Linguistics.
- Lorenzo Malandri, Fabio Mercurio, Mario Mezzanzanica, and Navid Nobani. 2021. [Taxoref: Embeddings evaluation for AI-driven taxonomy refinement](#). In *Machine Learning and Knowledge Discovery in Databases. Research Track - European Conference, ECML PKDD 2021, Bilbao, Spain, September 13-17, 2021, Proceedings, Part III*, volume 12977 of *Lecture Notes in Computer Science*, pages 612–627. Springer.
- Ken McRae, George S Cree, Mark S Seidenberg, and Chris McNorgan. 2005. Semantic feature production norms for a large set of living and nonliving things. *Behavior research methods*, 37(4):547–559.
- Kanishka Misra, Julia Taylor Rayz, and Allyson Ettinger. 2022. [A property induction framework for neural language models](#). *CoRR*, abs/2205.06910.
- Weiran Pan, Wei Wei, and Feida Zhu. 2022. [Automatic noisy label correction for fine-grained entity typing](#). In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 4317–4323. ijcai.org.
- Jeffrey Pennington, Richard Socher, and Christopher 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, Doha, Qatar. Association for Computational Linguistics.
- Julian Risch, Philipp Hager, and Ralf Krestel. 2021. [Multifaceted domain-specific document embeddings](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Demonstrations*, pages 78–83, Online. Association for Computational Linguistics.
- Alex Rosenfeld and Katrin Erk. 2023. [An analysis of property inference methods](#). *Nat. Lang. Eng.*, 29(2):201–227.
- Sascha Rothe and Hinrich Schütze. 2016. [Word embedding calculus in meaningful ultradense subspaces](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 512–517, Berlin, Germany. Association for Computational Linguistics.
- Jingchuan Shi, Jiaoyan Chen, Hang Dong, Ishita Khan, Lizzie Liang, Qunzhi Zhou, Zhe Wu, and Ian Horrocks. 2023. [Subsumption prediction for e-commerce taxonomies](#). In *The Semantic Web - 20th International Conference, ESWC 2023, Hersonissos, Crete, Greece, May 28 - June 1, 2023, Proceedings*, volume 13870 of *Lecture Notes in Computer Science*, pages 244–261. Springer.
- Nikhita Vedula, Patrick K. Nicholson, Deepak Ajwani, Sourav Dutta, Alessandra Sala, and Srinivasan Parthasarathy. 2018. [Enriching taxonomies with functional domain knowledge](#). In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR 2018, Ann Arbor, MI, USA, July 08-12, 2018*, pages 745–754. ACM.
- Ivan Vulić, Edoardo Maria Ponti, Anna Korhonen, and Goran Glavaš. 2021. [LexFit: Lexical fine-tuning of pretrained language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5269–5283, Online. Association for Computational Linguistics.
- Ivan Vulić, Edoardo Maria Ponti, Robert Litschko, Goran Glavaš, and Anna Korhonen. 2020. [Probing pretrained language models for lexical semantics](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7222–7240, Online. Association for Computational Linguistics.
- Weiqi Wang, Tianqing Fang, Wenxuan Ding, Baixuan Xu, Xin Liu, Yangqiu Song, and Antoine Bosselut. 2023a. [CAR: Conceptualization-augmented reasoner for zero-shot commonsense question answering](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13520–13545, Singapore. Association for Computational Linguistics.
- Weiqi Wang, Tianqing Fang, Baixuan Xu, Chun Yi Louis Bo, Yangqiu Song, and Lei Chen. 2023b. [CAT: A contextualized conceptualization and instantiation framework for commonsense reasoning](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13111–13140, Toronto, Canada. Association for Computational Linguistics.
- Chen Xing, Negar Rostamzadeh, Boris N. Oreshkin, and Pedro O. Pinheiro. 2019. Adaptive cross-modal few-shot learning. In *Proceedings of the Annual Conference on Neural Information Processing Systems*, pages 4848–4858.

Wenhan Xiong, Jiawei Wu, Deren Lei, Mo Yu, Shiyu Chang, Xiaoxiao Guo, and William Yang Wang. 2019. [Imposing label-relational inductive bias for extremely fine-grained entity typing](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 773–784, Minneapolis, Minnesota. Association for Computational Linguistics.

Kun Yan, Zied Bouraoui, Ping Wang, Shoaib Jameel, and Steven Schockaert. 2021. [Aligning visual prototypes with BERT embeddings for few-shot learning](#). In *Proceedings of the International Conference on Multimedia Retrieval*, pages 367–375.

A Training Details

Predicting Commonsense Properties We use the AdamW optimizer with learning rate $2e-5$. We use early stopping with a patience of 20. We use the same settings as [Gajbhiye et al. \(2022\)](#) for both fine-tuning and pre-training. During pre-training, we randomly select 10000 pairs for tuning. For the concept split we use the fixed training-test split from [Gajbhiye et al. \(2022\)](#). For property split, we use 5-fold cross-validation, while for C+P, we used the 3×3 fold cross-validation strategy. During fine-tuning, for all splits, we randomly select 20% from the training set as validation data for model selection. We use a batch size of 32 and set the maximal sequence length for the concept and property prompts to 32.

Outlier Detection For these experiments, we use the same bi-encoder models as for predicting commonsense properties. Specifically, we have used the model initialised from BERT-large-uncased, which was trained using ConceptNet and ChatGPT. However, in this case, the model is used without further fine-tuning.

Ontology Completion To get concept clusters from the affinity propagation algorithm, we tune the preference values from $\{0.5, 0.6, 0.7, 0.8, 0.9\}$. We set the learning rate to $1e-2$, the maximum number of epochs to 200 and the dropout rate to 0.5. We use AdamW as optimizer and tune the number of hidden dimensions from $\{8, 16, 32, 64\}$ and the number of GNN layers from $\{3, 4, 5\}$. We use a weight decay of $5e-2$.

Ultra-Fine Entity Typing To get concept clusters, we again use affinity propagation and select the preference values from $\{0.5, 0.6, 0.7, 0.8, 0.9\}$. We use a learning rate of $2e-5$ with AdamW as op-

timizer, a batch size of 16 and a maximum number of 500 epochs.

B Outlier Detection Dataset

Table 5 shows the properties that were selected from the McRae dataset for constructing the outlier detection benchmark. For each property, we selected concepts that are asserted to have this property in the dataset, and we organised them into broad taxonomic categories, which are also shown in Table 5.

C Qualitative Analysis

Table 6 shows examples of the nearest neighbours of *frisbee* and *bureau* in some of the facet-specific embedding spaces that are used by the MClu strategy. For this analysis, we use the set of concepts from the McRae dataset. These examples illustrate how different facets emphasise different aspects.

For the case of *frisbee*, the first facet links this concept to other sports related terms. In the second facet, it is instead related to round things. In the third facet, the top neighbours are related to kids. Due to the way in which the facets are learned (e.g. by considering a fixed number of facet embedding clusters), there are also some facets that reflect a mixture of different aspects. For instance, the last facet in Table 6 combines elements of the first three facets, i.e. the nearest neighbours cover sports related concepts, kids related concepts, and round things. For *bureau*, in the first facet we find office related terms. In the second facet, the top neighbours are different types of furniture. In facets 3 and 4 we see a mixture of different kinds of terms.

Table 7 similarly shows examples of nearest neighbours in different facets for the label space from the UFET dataset.

Property	Positive Examples
Dangerous	1: alligator, bear, beehive, bull, cheetah, crocodile, lion, rattlesnake, tiger 2: bazooka, bomb, bullet, crossbow, dagger, grenade, gun, harpoon, missile, pistol, revolver, rifle, rocket, shotgun, sword, axe, baseball bat, knife, machete 3: motorcycle
Edible	1: apple, banana, cranberry, tomato, tangerine, strawberry, spinach, rhubarb, raspberry, radish, pumpkin, prune, mushroom, parsley, walnut, rice, raisin, potato, plum, pineapple, pepper, peas, pear, peach, orange, onions, olive, nectarine, lime, lettuce, lemon, grapefruit, grape, garlic, cucumber, corn, coconut, cherry, celery, cauliflower, carrot, cantaloupe, cabbage, broccoli, blueberry, beets, beans, avocado, asparagus 2: bread, cake, cheese, hot dog, pizza, sandwich, pie, donut, biscuit 3: crab, deer, hare, octopus, salmon, turkey, tuna, trout, squid, shrimp, sardine, pig, octopus, lobster, lamb, goat, cow, clam, chicken
Flies	1: bird, crow, dove, duck, eagle, falcon, flamingo, goose, hawk, woodpecker, pigeon, owl, peacock, seagull 2: butterfly, hornet, housefly, wasp, moth 3: airplane, helicopter, jet, missile, rocket 4: balloon, kite, frisbee
Hot	1: bathtub 2: cigar, cigarette, candle 3: hair drier, kettle, oven, stove, toaster
Lives in water	1: alligator, crocodile, otter, turtle, seal, frog, flamingo, salamander, swan, walrus 2: clam, crab, lobster, octopus, squid, shrimp 3: dolphin, eel, goldfish, whale, salmon, sardine, trout, tuna 4: boat, sailboat, ship, submarine, yacht, canoe
Produces noise	1: accordion, bagpipe, clarinet, flute, harp, piano, trombone, violin 2: airplane, ambulance, helicopter 3: bomb, cannon, grenade 4: cell phone, hair drier, stereo
Sharp	1: axe, bayonet, dagger, machete, knife, spear, sword 2: chisel, corkscrew, screwdriver, scissors 3: blender, grater 4: razor 5: pin
Used by children	1: balloon 2: buggy 3: crayon, paintbrush, pencil 4: doll, teddy bear, toy 5: earmuffs 6: frisbee, kite 7: skateboard, sled, tricycle
Used for cooking	1: apron 2: blender, grater, mixer 3: bowl, colander, pan, pot, skillet, strainer 4: kettle, microwave, oven, stove, toaster 5: knife, ladle, spatula, spoon, tongs
Worn on feet	1: boots, sandals, shoes, slippers 2: nylons, socks 3: skis, snowboard

Table 5: Results for outlier detection in terms of F1 (%).

Concept	Neighbours
frisbee	<i>Facet 1:</i> tricycle, surfboard, sports_ball, tennis_racket, snowboard, kite, doll, balloon, toy, pie
	<i>Facet 2:</i> balloon, pie, sports_ball, cake, donut, teddy_bear, surfboard, kite, doll, toy
	<i>Facet 3:</i> kite, doll, balloon, toy, tricycle, moth, pie, sports_ball, surfboard, football
	<i>Facet 4:</i> tricycle, surfboard, sports_ball, tennis_racket, snowboard, kite, doll, balloon, toy, pie
bureau	<i>Facet 1:</i> envelope, certificate, typewriter, doorknob, fence, cabinet, carpet, shelves, bookcase, gopher
	<i>Facet 2:</i> desk, dining_table, table, shelves, bookcase, envelope, typewriter, gopher, escalator, peg
	<i>Facet 3:</i> shelves, bookcase, envelope, desk, peg, gopher, tack, cabinet, handbag, hook
	<i>Facet 4:</i> bookcase, cabinet, shelves, desk, doorknob, dining_table, envelope, typewriter, handbag, lamp

Table 6: Nearest neighbours for different facets, using the concepts from the McRae dataset.

Concept	Neighbours
keyboard	<i>Facet 1:</i> bass, rhythm guitar, lead guitar, drum major, pianist, bass drum, bass guitar, bassist, bass guitarist, air guitar
	<i>Facet 2:</i> processor, organ, desktop, mac, storage, file system, thumb drive, file extension, touch screen, desktop environment
business card	<i>Facet 1:</i> flash card, smart card, bank card, graphics card, green card, index card, network card, phone card, press card, punch card, report card, trade card
	<i>Facet 2:</i> business class, business cycle, business day, business economics, business end, business ethics, business intelligence, business logic

Table 7: Nearest neighbours for different facets, using the labels from the UFET dataset.