


# On Entity Identification in Language Models

Masaki Sakata<sup>1,2</sup> Benjamin Heinzerling<sup>2,1</sup> Sho Yokoi<sup>3,1,2</sup>  
Takumi Ito<sup>1,4</sup> Kentaro Inui<sup>5,1,2</sup>

<sup>1</sup> Tohoku University <sup>2</sup> RIKEN <sup>3</sup> NINJAL <sup>4</sup> Langsmith Inc. <sup>5</sup> MBZUAI  
sakata.masaki.s5@dc.tohoku.ac.jp yokoi@ninjal.ac.jp  
benjamin.heinzerling@riken.jp t-ito@tohoku.ac.jp kentaro.inui@mbzuai.ac.ae

## Abstract

We analyze the extent to which internal representations of language models (LMs) identify and distinguish mentions of named entities, focusing on the many-to-many correspondence between entities and their mentions. We first formulate two problems of entity mentions — ambiguity and variability — and propose a framework analogous to clustering quality metrics. Specifically, we quantify through cluster analysis of LM internal representations the extent to which mentions of the same entity cluster together and mentions of different entities remain separated. Our experiments examine five Transformer-based autoregressive models, showing that they effectively identify and distinguish entities with metrics analogous to precision and recall ranging from 0.66 to 0.9. Further analysis reveals that entity-related information is compactly represented in a low-dimensional linear subspace at early LM layers. Additionally, we clarify how the characteristics of entity representations influence word prediction performance. These findings are interpreted through the lens of isomorphism between LM representations and entity-centric knowledge structures in the real world, providing insights into how LMs internally organize and use entity information. : <https://github.com/masaki-sakata/entity-identification>

## 1 Introduction

Transformer-based language models (LMs) have been demonstrated to be able to recall factual knowledge composed of entities and relations (Petroni et al., 2019; Jiang et al., 2020; Heinzlerling and Inui, 2021; Cohen et al., 2023). For example, when given the query “Barack Obama was born in \_\_\_”, LMs can predict appropriate words like ‘Hawaii’. To understand how LM predicts such entities (e.g. “Hawaii”), extensive internal analyzes have been performed from an informa-

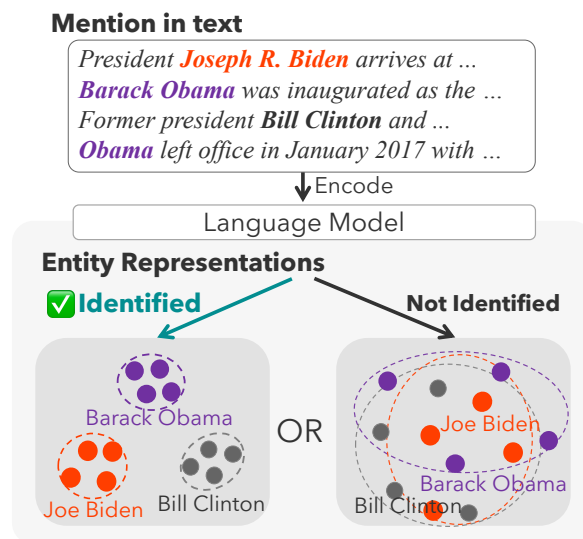


Figure 1: Illustration of *entity identification*. If “Barack Obama” or “Joe Biden” are represented in various mentions and contexts but still form a cluster as a single entity in the representation space, then we observe successful entity identification. Such entity identification in the representation space suggests that the LM’s representations reflect certain aspects of the structure of real-world, entity-centric knowledge.

tion flow perspective (Dai et al., 2022; Meng et al., 2022; Geva et al., 2023).

However, while previous work has shed light on the internal mechanisms of LMs (Dai et al., 2022; Meng et al., 2022; Geva et al., 2023), our understanding remains limited in scope. For instance, when processing mentions like “Obama” or “Barack Obama”, it remains unclear whether LMs encode them as the same individual based on context (Figure 1). Although encoding mentions as distinct entities may appear straightforward, there are two major real-world challenges (Figure 2). One is **mention ambiguity**: a single mention could potentially refer to multiple different entities. The other is **mention variability**: multiple mentions with different surface forms exist for a single entity. No-

tably, previous studies have not addressed queries containing mention ambiguity or mention variability, leaving a gap in our understanding of how LMs handle these fundamental challenges. Properly resolving mention ambiguity and variability is crucial to language understanding.

We assume that encoding where mentions are distinguished by their corresponding entities, as shown in Figure 1 (left), is an appropriate representation and investigate to what extent LMs distinguish entities in their internal representations. We call the process of distinguishing entities in internal representations “entity identification”. To evaluate entity identification, we used Purity (Zhao and Karypis, 2001) and Inverse Purity (IP). This approach allows us to directly measure whether the model’s representations for the same entity form compact clusters that differ from other entities.

Our experiments with five autoregressive LMs demonstrated that they achieve an AUC of approximately 0.8–0.9 for mention ambiguity (§4.1) and 0.66–0.8 even in the presence of mention variability (§4.2). LM scores outperformed the baselines, demonstrating their ability to identify entities. We also find that entity information forms low-dimensional subspaces in early layers (§5.1), with better entity representation separation improving word prediction (§5.2). Finally, we discuss these findings through the lens of isomorphism between LM representations and real-world entity-centric knowledge structures (§6). This provides evidence that LMs encode discrete structural knowledge through text-only training, extending previous observations (Abdou et al., 2021; Chen et al., 2023; Gurnee and Tegmark, 2024; Park et al., 2024b) toward a systematic understanding of how LM representation spaces correspond to structural knowledge.

## 2 Entity Identification

Simply put, our goal is to find out if LMs “know”<sup>1</sup> that different mentions of the same real-world entity refer to the same entity and, conversely, that mentions of different entities refer to different entities. For example, in discussions about “Barack Obama” or “President Obama”, a LM with world knowledge should be able to recognize that both mentions refer to the same person.

<sup>1</sup>This loose, anthropomorphizing use of the word “know” will be made more precise in the following section.

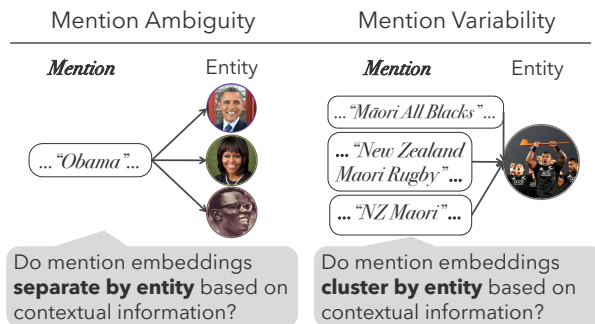


Figure 2: Two major factors that can make entity identification difficult are ambiguity and variability of entity mentions.

### 2.1 Why is Entity Identification Difficult?

What makes entity identification difficult can be attributed to the ambiguity and variability of a mention (Figure 2). From the perspective of mention ambiguity, “Obama” could refer to “Barack Obama” or “Michelle Obama” depending on the context. From the perspective of mention variability, for example, “Māori All Blacks” and “New Zealand Maori Rugby” refer to the same rugby team in New Zealand despite using different words. Resolving mention ambiguity and variability requires effective use of surface and contextual information.

### 2.2 Purity and Inverse Purity

When a mention with context is input, we aim to quantify the degree to which the LM distinguishes between the entities corresponding to that mention in its representations. Intuitively, if the LM effectively distinguishes between entities, the embeddings representing these entities are expected to be clearly separated based on conceptual similarity (as illustrated in the example of “Barack Obama” in Figure 1). To measure how well the embeddings of mentions are separated and encoded on an entity basis, we use an F1 score composed of (local) **Purity** (Zhao and Karypis, 2001) and (local) **Inverse Purity (IP)**.<sup>2</sup> In other words, we assess the geometric locality of a set of embeddings by examining how much they are mixed with other sets of embeddings. In Figure 3, we illustrate several embeddings, showing both the **class** division cor-

<sup>2</sup>As an alternative to Purity and IP, the Adjusted Rand Index (ARI) (Hubert and Arabie, 1985) can also be employed. ARI was implemented in this study, and the results are reported in Appendix D; they show the same overall trend as the Purity/IP scores. Although ARI corrects for chance agreement, it is less suitable for fine-grained, entity-level analysis and is therefore omitted from the main text.

responding to the Biden entity (**Biden Class**) and the **cluster** division formed based on the centroid of the Biden embeddings (**Biden Cluster**). The F1 score using (local) Purity and (local) IP indicates how well the **Biden Cluster** matches the true **Biden Class**. An F1 score of 1.0 indicates that the Biden embeddings are completely separated from other clusters. As the score decreases, it indicates that the embeddings are increasingly mixed with other clusters. The detailed computational procedure for Purity, IP and F1 score can be found in the Appendix A. Note that since this evaluation method relies on distances between representations, it may be affected by the curse of dimensionality, which we examine in § 3.4.

The F1 score based on purity and IP can be considered as one of the methodological approaches to measure how well LMs capture the structure of the world (see Table 1 in Appendix). There is widespread interest in understanding how well LMs capture the structure of the world. Typically, the *degree* of correspondence between the world’s inherent structure (such as the similarity between colors in color space) and their representation within LMs (such as the similarity between color names in the LM’s latent space) is evaluated through methods like representational similarity analysis (RSA) (Abdou et al., 2021; Patel and Pavlick, 2022; Chen et al., 2023; Gurnee and Tegmark, 2024; Hernandez et al., 2024; Park et al., 2024b). For entities, however, the crucial aspect is whether pairs of mentions refer to the same entity or different ones. Our proposed clustering-based method measures the *degree* of alignment between this binary structure (“same or different”) in the world and in LMs. This approach enables a more natural evaluation of word pairs that, while semantically very similar, would compromise factuality if substituted for one another (such as mentions of entities).

### 3 Experimental Setup

#### 3.1 Models

We used five variants of LMs, including GPT-2 (Radford et al., 2019), Llama-2 (7B, 13B) (Touvron et al., 2023), Llama-3 (8B) (Dubey et al., 2024), and Mistral (7B) (Jiang et al., 2023). The details are shown in Appendix B.1, Table 3.

We use the following three baselines: (1) Random embeddings: Individual embeddings randomly sampled from a normal distribution for each mention occurrence. Different embeddings are as-




Figure 3: Overview of local Purity and local IP based on Equation 4 and Equation 6. As in Equation 8, a high F1 score (purity and IP) indicates not only that embeddings of the same **class** are contained within a **cluster**, but also that embeddings of each **class** are appropriately separated into distinct **cluster**. The circles represent the embeddings corresponding to each mention.

signed even for identical mentions. (2) Unique mention embeddings: Fixed embeddings randomly sampled from a normal distribution for each unique mention. The same embedding is always assigned to identical mentions (e.g., the mention “Obama” consistently uses the same embedding). (3) Fast-Text embeddings<sup>3</sup> (Mikolov et al., 2018): Pre-trained static word embeddings. When a mention consists of multiple tokens, we use the average embedding of those tokens. Unlike LMs, unique mention embeddings and FastText embeddings only capture surface-level information of mentions.

#### 3.2 Entity Representations

Following the configurations of Abdou et al. (2021); Bommasani et al. (2020); Vulić et al. (2020), we use the hidden states of LMs as representations of a given word. The term “Entity Representations” refers to the representations that correspond to the mention of an entity. In the case of autoregressive LMs, several works have indicated that information related to an entity is concentrated in the last token of a mention (Meng et al., 2022; Geva et al., 2023; Heinzerling and Inui, 2024). Thus, we report the results using embeddings of the last token.<sup>4</sup> Since autoregressive LMs cannot read the context that appears after analyzing

<sup>3</sup>We use embeddings trained on Wikipedia 2017, the UMBC webbase corpus (Han et al., 2013), and the statmt.org news dataset.

<sup>4</sup>The results using average embeddings of subwords are reported in Appendix C.2.2 and C.3.2.




Figure 4: Comparison of F1 scores across different dimensions for LMs, FastText, and random embeddings. The vertical dashed line indicates dimension=20. For each model, the layer with the highest AUC score is visualized.

the mention, we repeat the input sentence following Springer et al. (2024). Specifically, the autoregressive LM reads the entire sentence once and then uses the embedding of the mention appearing the second time as the entity representation.<sup>5</sup>

### 3.3 Data

We used the ZELDA-TRAIN dataset (Milich and Akbik, 2023), employed for the entity disambiguation task. We filter out entities with fewer than five instances, as sparse instances could lead to biased centroid positions when computing the average semantic representation of each entity. Additionally, instances where the length of the tokenized sentences exceeded 500 tokens were excluded to ensure compatibility with a wide range of LMs, setting a conservative input length limit. Table 5 in Appendix B.3 shows the statistical information of the dataset.

### 3.4 Preliminary Study: Dimensional Impact

We first investigate the effect of dimensionality on entity identification. This is crucial because our evaluation method focuses on the distance between representations, making it potentially susceptible to the curse of dimensionality. Using the filtered ZELDA-TRAIN dataset, we compared three kinds of representations: (i) LMs, (ii) FastText, and (iii) random embeddings. We reduced the LM and FastText representations with linear discriminant analysis (LDA). Random embeddings, in contrast, were generated directly at each target dimensionality. Figure 4 summarizes the results and reveals two main observations. First, random embeddings score almost zero in low dimensions but nearly saturate in high dimensions. This indicates that

<sup>5</sup>Given a sentence, “Alice went to Paris,” we input “Alice went to Paris. Alice went to Paris” to the LM and used the embedding of the second “Alice” as the entity representation.

our metric is affected by the curse of dimensionality. Second, LMs remain strong even in low dimensions; for example, Llama-2-7B attains an F1 of 0.90 in 20 dimensions, only 3% lower than its score in 4,096 dimensions. To determine the optimal dimension, we analyze the rate of performance change ( $\Delta F1/\Delta d$ ) in different dimension windows. The slope of the F1 score of LMs rapidly decreases to approximately 20 dimensions, after which it approaches zero ( $|\Delta F1/\Delta d| < \epsilon$  where  $\epsilon = 0.005$  for  $d > 20$ ). Therefore, we conduct subsequent analyses in 20-dimensional space where the curse of dimensionality has minimal impact. This efficient encoding of entity information in low-dimensional subspaces aligns with the manifold hypothesis of LM representations (Cai et al., 2021; Cheng et al., 2023; Valeriani et al., 2023).

## 3.5 Definition of Difficulty

### 3.5.1 Mention Ambiguity

Mention ambiguity quantifies how ambiguously a mention is assigned to different candidate entities based on its frequency distribution between entities. For example, the mention “Oxford” can refer to different entities depending on the context, such as Oxford University or Oxford as a place name, and these entities are referenced with similar frequency. When a mention refers to multiple entities with similar probabilities, the ambiguity becomes high in such cases. We use the entropy  $H$  of the candidate mappings from mention to entity as the mention ambiguity:  $H = -\sum_{i=1}^N p_i \log p_i$ . Here,  $p_i$  represents the proportion of times a mention refers to the candidate entity  $i$ .<sup>6</sup> A higher entropy  $H$  indicates greater uncertainty in the candidate entity corresponding to a mention, suggesting more ambiguity (see Table 2 in Appendix). In our experiments, we used a subset of ambiguous mentions, excluding cases where  $H = 0$ .

### 3.5.2 Mention Variability

The variability of the mention refers to the degree to which different surface forms can vary when referring to the same entity. For example, the mentions “Māori All Blacks” and “New Zealand Maori Rugby” refer to the same entity, but their surface forms are significantly different. In such cases, we want to assign a high score to mention the variability. We use surface-form dissimilarity, which indicates the surface-level variation of multiple men-

<sup>6</sup>The detailed definition of mention ambiguity is provided in Appendix B.4.



tions for the same entity, as mention variability. This dissimilarity  $D$  is formulated as:

$$D = \frac{2}{|M_e|(|M_e| - 1)} \sum_{i=1}^{|M_e|-1} \sum_{j=i+1}^{|M_e|} \frac{L(m_i, m_j)}{\max(|m_i|, |m_j|)}$$

$L(m_i, m_j)$  is the Levenshtein distance<sup>7</sup> between mentions  $m_i$  and  $m_j$ , and  $|M_e|$  is the total number of mentions for the same entity.<sup>8</sup> A higher dissimilarity  $D$  indicates greater surface-level differences between mentions for the same entity (see Table 2 in Appendix).

## 4 Results

### 4.1 Mention Ambiguity

In the upper-left region of Figure 5 shows the AUC scores of the baseline model and the highest AUC scores achieved across all layers of each LM when controlling for mention ambiguity. Higher scores on the X-axis indicate more ambiguous mentions, while higher scores on the Y-axis indicate more successful entity identification (entity separated).

LMs achieved AUC scores of 0.8–0.9. This score is markedly higher than the AUC of baselines.<sup>9</sup> These results demonstrate that LMs distinguish the identity of entities even for ambiguous mentions, compared to baselines.

We found that LMs achieve effective entity identification by stacking layers and contextualizing embeddings. As shown in the bottom left of Figure 5, Llama-2 7B’s AUC substantially improved from 0.38 at layer 0 to 0.87 at layer 8. This transformation can also be qualitatively observed in Appendix Figure 10. At layer 0, embeddings of the country name “Georgia” and the U.S. state name “Georgia” are mixed, but by layer 8, they are clearly separated. Thus, contextual information is essential for entity identification in cases with mention ambiguity, and LMs can address this challenge through embedding contextualization by stacking layers.

<sup>7</sup>Levenshtein distance is calculated at the character level.

<sup>8</sup>The detailed definition of mention variability is provided in Appendix B.5.

<sup>9</sup>Random Embeddings generate distinct embeddings for each instance of the same mention, potentially leading to spuriously inflated scores in mention ambiguity. For instance, when the mention “Obama” occurs multiple times in the text, each occurrence is assigned a different embedding. While this might result in an embedding space that appears to differentiate between Barack Obama and Michelle Obama, such discrimination is purely coincidental.

### 4.2 Mention Variability

In the upper-right region of Figure 5 shows the AUC scores of the baseline model and the highest AUC scores achieved across all layers of each LM when controlling for mention variability. Higher scores on the X-axis indicate more variable mentions, while higher scores on the Y-axis indicate more successful entity identification. The results showed similar trends to mention ambiguity.<sup>10</sup>

Overall, LMs achieved AUC scores of 0.66–0.8. This score stands in stark contrast to the AUC of random embeddings, unique mention embeddings, and FastText. These results demonstrate that LMs distinguish entity identities even with variable mentions, in contrast to the baselines. However, all LMs showed a notable decrease in F1 scores when mention variability exceeded 0.8. For example, the basketball team Sasaki Baskonia, also known as “Taugrés,” is a case with a variability exceeding 0.8. In such cases, Llama-2 7B recorded an F1 score of 0.5.

We found that LMs achieve effective entity identification by stacking layers and contextualizing embeddings. As shown in the bottom right of Figure 5, Llama-2 7B’s AUC substantially improved from 0.3 at layer 0 to 0.81 at layer 8. Thus, contextual information is essential for entity identification in cases with mention variability, and it was confirmed that LMs can effectively address this challenge through embedding contextualization by stacking layers. However, it’s notable that AUC scores decrease in the later layers (16 and 32) compared to layer 8. One possible explanation is that the influence of next-word prediction hinders the process of distinguishing entities in the input representations.

## 5 Additional Experiments

We conduct additional experiments to analyze two points: (1) the properties of entity representations (e.g., linear separability) and (2) their impact on word prediction.

### 5.1 Properties of Entity Representations

**Entity Representations are Linearly Separable in Early Layers** To determine whether entity representations are encoded in a linearly separable

<sup>10</sup>F1 scores cannot be directly compared between mention ambiguity and mention variability due to different evaluation data units. Mention ambiguity is evaluated on data aggregated by identical mentions, while mention variability is evaluated on an entity basis.

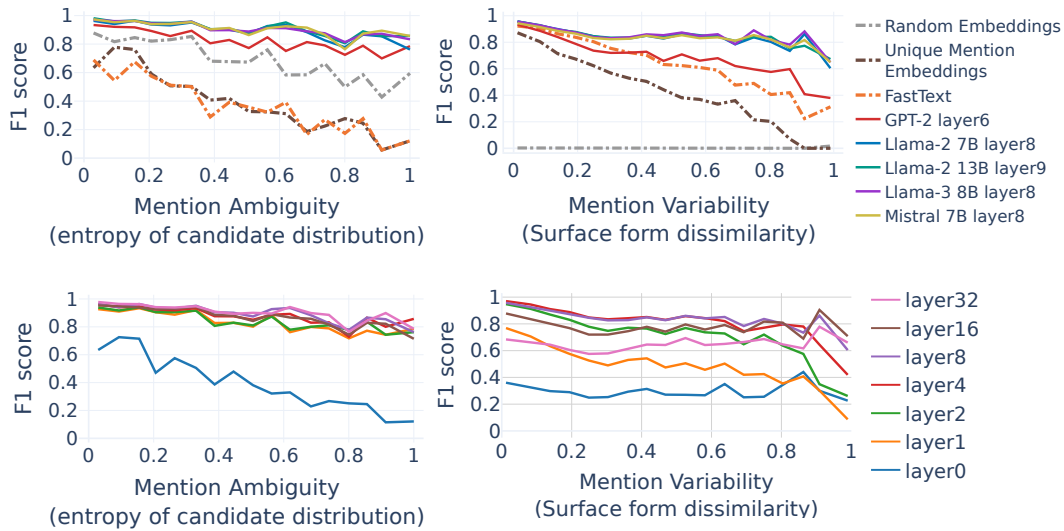


Figure 5: Separability (F1 scores) of entity representations for each difficulty (left: mention ambiguity, right: mention variability). F1 scores are averaged within each bin on the x-axis. Top: Results across all models, using the layer with highest AUC for LM scores. Bottom: Layer-wise analysis for Llama-2 7B. Complete results and AUC scores for all models are provided in Appendix C.2.1 and C.3.1.

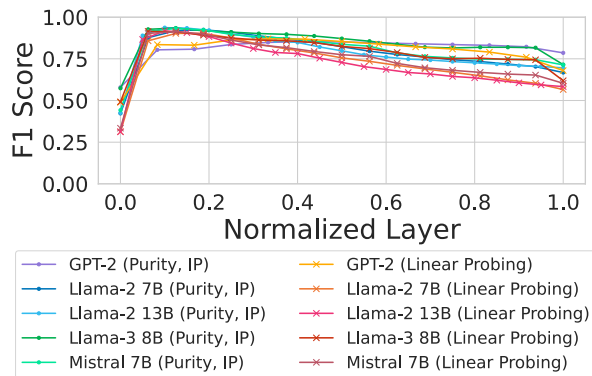


Figure 6: F1 scores based on Purity, IP, and Linear Probing. The x-axis represents the relative position of the layers. All embeddings are reduced to 20 dimensions.

manner, we conduct linear probing. Linear probing serves as an alternative method to Purity and IP, involving training a linear classifier to predict entity labels from representations. The detailed settings for linear probing are described in Appendix B.2. To investigate which layers contain information necessary for entity identification and whether there are differences in trends among Purity, IP, and Linear Probing results, we calculate F1 scores for each layer. As in §4, the embedding dimensions are reduced to 20 dimensions using LDA. The results of these analyses are shown in Figure 6.

Our analysis revealed two main findings. First, Linear Probing achieved F1 scores of around 0.9. This suggests that entity representations are encoded in an (almost) linearly separable form. Previ-

ous studies (Olah et al., 2020; Elhage et al., 2022; Gurnee and Tegmark, 2024) have been accumulating evidence supporting the linear representation hypothesis, which suggests that features within neural networks are represented linearly. Our results is further evidence to support this hypothesis.

Second, we found that the maximum F1 score peaks for both Purity, IP, and Linear Probing occur around normalized layer 0.2 (e.g., layers 6-8 in Llama-2 7B). This indicates that the information necessary for entity identification is present in the early layers. These layer-wise linear probing results align with the findings of Gurnee and Tegmark (2024), who reported that probing scores for geographical and temporal features peak in early layers. While they focused specifically on geographical and temporal features, our analysis extends to a broader range of entities. Our results also show similarities with Meng et al. (2022), who found that entity information is encoded in early layers when processing the last token of entity mentions. They showed this using activation patching, while we confirmed the same trend through a different approach that examines geometric properties in embedding space. Considering that our experiments were conducted in a 20-dimensional space, these results suggest that the information necessary for entity identification is encoded in a low-dimensional linear subspace of the early layers in LMs.



Figure 7: Results of RSA experiments. Each value represents Spearman’s rank correlation coefficient. LM representations are taken from the layer with the highest F1 score. Dimensionality is fixed at 20 for all representations.

### Structural Similarity of Entity Representations

Examining the upper portion of Figure 5, we observe similar F1 score patterns across different LMs, particularly among Llama-2, Llama-3, and Mistral. This pattern suggests that these models develop comparable structural representations for entities. To quantify this observation, we conducted Representation Similarity Analysis (RSA) (Kriegeskorte et al., 2008) to measure the structural similarity between representations (see Appendix B.7 for experimental details).

Figure 7 presents our results. First, we observe that pairs of LM representations exhibit higher similarity than the static embedding baseline. Interestingly, Llama-2, Llama-3, and Mistral exhibit stronger inter-model similarity than comparisons involving GPT-2. This finding suggests that larger, more capable LMs converge toward similar entity representation structures. While our results echo the findings of Huh et al. (2024), we extend their work by focusing specifically on entity representations. This structural similarity allows us to interpret LMs through the lens of isomorphism, which we discuss in detail in § 6.

### 5.2 Effect of Entity Representations Structure on Word Prediction

To understand how the degree of entity representation separation affects word output, we investigate from two perspectives: output consistency and entity disambiguation accuracy.

**Output Consistency** Cao et al. (2022) reported cases where simply changing “The U.S.” to “Amer-

ica” in the prompt “The U.S. capital of \_\_\_” led to different word prediction results. In other words, they reported that the variation in surface forms used to refer to the same entity causes inconsistency in LM’s word predictions. In such cases, it would be desirable for the model to consistently predict “Washington” for both inputs. We verify through a simple experiment that this word prediction inconsistency is related to the confusion in entity representations. In the experiment, we used prompts such as “[X] was founded in \_”. We compared the consistency of model responses by substituting mentions of two types of entities with different mention variability for [X]. As discussed in § 4.2, entities with low mention variability tend to have mentions encoded as a clustered representation, while entities with high mention variability show confusion between representations. The intuition behind this experimental setup is that differences in mention variability cause structural differences in representations, which affect variations in word prediction. We sampled 20 organization entities each with low (0-0.3) and high (0.7-1.0) mention variability and predicted [Object] using the format “[Subject] [Relation]” with mentions belonging to these entities as Subject.<sup>11</sup> Word prediction consistency was calculated by collecting words for [Object] (the first word formed by combining generated tokens) and computing the proportion of the most frequent prediction within the same entity. Then, we calculated the average for both low and high mention variability groups.

The results revealed that while word prediction consistency averaged 71% for low mention variability, it significantly decreased to an average of 39% for high mention variability. In other words, we confirmed that higher mention variability leads to lower model prediction consistency, supporting the hypothesis that the internal state of entity representations confusion causes prediction inconsistency. A more detailed qualitative analysis of cases with inconsistent word prediction is presented in Appendix C.4.

### Entity Disambiguation Accuracy

We investigate how the degree of separation in entity representations affects entity disambiguation accuracy. Intuitively, if entity representations are well-separated, the model should successfully perform entity dis-

<sup>11</sup>For Relations, we used “[X] was founded in [Y].” and “The headquarter of [X] is in [Y].” from LAMA Probe (Petroni et al., 2019). Llama-2 7B was used as the model.

ambiguity. To examine the relationship between representation separation and entity disambiguation accuracy, we conducted in-context learning experiments as follows. We input the following prompt to the LM: Does **X** refer to A or B? One-word answer only. A: San Diego B: San Francisco Answer: . We then use PatchScope (Ghandeharioun et al., 2024) to patch contextualized mention embeddings of San Diego or San Francisco into the position marked as **X**. If the entity representations of San Diego and San Francisco are well-separated, the model should correctly resolve the entity disambiguation task. We used approximately 7,800 entities, consistent with the dataset size in § 4. We employed 10-shot learning and constrained the model to output only “A” or “B” as answers. We selected Llama-2 13B as it supports few-shot learning. Detailed experimental settings are provided in Appendix B.8.

Figure 8 presents our results. Figure 8-(a) shows the effect of varying the layer used for patching. When using representations from early layers, entity disambiguation accuracy reaches approximately 74%, significantly exceeding the 50% chance rate. This finding Accuracy aligns with the trends observed in § 5.1, suggesting that information from early layers contributes to entity disambiguation performance. Figure 8-(b) plots entity disambiguation accuracy against F1 score using layer 4, which achieved the highest entity disambiguation accuracy. The Pearson correlation coefficient is approximately 0.18, with a regression coefficient  $\beta$  of 0.31. The data shows an upward trend, indicating that entities with well-separated representations tend to achieve higher entity disambiguation accuracy.

## 6 Discussion: Insights from Isomorphism

In § 4, we revealed the extent to which LMs can perform entity identification. Interpreting these experimental results from the perspective of isomorphism described in § 2.2 and Table 1 in the Appendix, we can conclude that LMs exhibit a higher degree of isomorphism between the structure of entity representations and the structure of entity classes compared to FastText and random baselines. LMs achieved an AUC of approximately 0.9 in cases with mention ambiguity (§ 4.1) and an AUC of about 0.8 in cases with mention variability (§ 4.2), both scores being significantly higher than baselines. Furthermore, we observed simi-

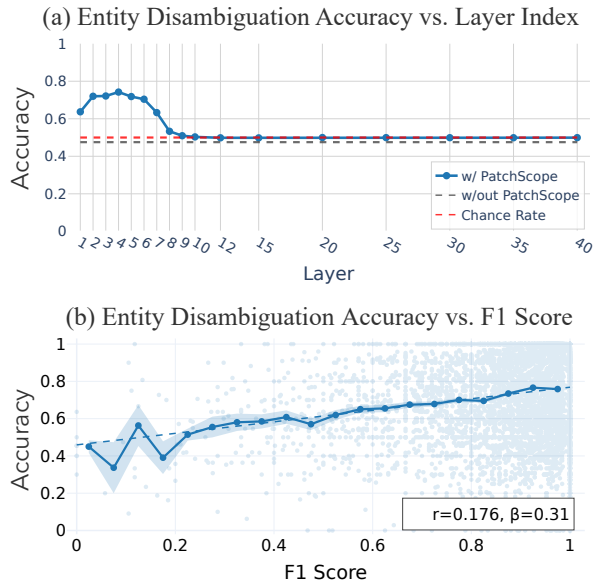


Figure 8: (a) Entity disambiguation accuracy for patch embeddings extracted from different Llama-2 13B layers. (b) Entity disambiguation accuracy vs. representation separation (F1 score) for 4th layer patch embeddings from Llama-2 13B. Entity disambiguation accuracies are averaged within each bin on the F1 score.  $r$ : Pearson correlation,  $\beta$ : regression coefficient.

larities in the structure of entity representations across different LMs (§ 5.1). These results experimentally support The Platonic Representation Hypothesis (Huh et al., 2024). Specifically, the similarity in entity representation structures across different LMs and their high isomorphism suggest that entity representation structures are “converging” toward a certain point, which can be interpreted as the structure of real-world entities (in this study, the discrete topology of “mentions belonging to the same entity class”). Our findings reinforce previous empirical observations that LMs can encode specific knowledge structures existing in the real-world through text-only learning (Abdou et al., 2021; Patel and Pavlick, 2022; Chen et al., 2023; Gurnee and Tegmark, 2024; Hernandez et al., 2024; Park et al., 2024b). Additionally, in § 5.2, we demonstrated that the structure of representations affects both the consistency of word predictions and entity disambiguation performance. We consider this one of the positive aspects of achieving high isomorphism.

## 7 Related Work

We review methods for verifying LMs knowledge and position our research by organizing existing



findings on knowledge structures.

### 7.1 Methods for Analyzing LM Knowledge

While there are numerous studies examining what knowledge LMs possess, their methodologies can be broadly categorized into three approaches. (i) **Behavioral analysis** (Petroni et al., 2019; Jiang et al., 2020; Marjeh et al., 2023; Patel and Pavlick, 2022) explores knowledge encoded in LMs using prompts designed to elicit responses. For example, evaluating whether the model can predict “Hawaii” from the prompt “Barack Obama was born in \_”. While this approach is interpretable to humans and provides intuitive insights, it has been reported that different phrasings of prompts with the same meaning can yield different prediction results (Jiang et al., 2020; Heinzlerling and Inui, 2021). Additionally, treating the model as a black box makes it difficult to understand how specific input words are encoded within the LM.

(ii) **Classifier-based probing** (Alain and Bengio, 2017; Hewitt and Manning, 2019; Liétard et al., 2021; Abdou et al., 2021; Gurnee and Tegmark, 2024) evaluates whether LMs’ representations encode specific information by measuring the performance of trained classifiers. These analysis results depend on classifier performance and can vary significantly based on optimizer and initialization settings (Zhou and Srikumar, 2021).

(iii) **Representational analysis** (Zhou and Srikumar, 2021; Abdou et al., 2021; Park et al., 2024b,a) analyzes LMs’ internal representations using unsupervised methods. For instance, when measuring the similarity between two representational structures, it compares distance matrices between embedding pairs. This approach can evaluate representations directly without classifiers, making it less susceptible to optimization effects and enabling a purer understanding of the model’s internal structure. The Purity (Zhao and Karypis, 2001) and IP metrics used in this study belong to this representational analysis category.

### 7.2 Structure in LM Representations

The correspondence between LMs’ representational spaces and real-world knowledge structures has been studied from various perspectives. Mikolov et al. (2013) demonstrated that word embeddings can linearly represent semantic relationships in vector space. Subsequently, Petroni et al. (2019) and Poerner et al. (2020) used behavioral analysis to prove that pre-trained LMs store factual

knowledge composed of entities and relations. For instance, they confirmed that LMs could predict appropriate words like “Hawaii” from inputs such as “Barack Obama was born in \_\_\_”. The observation that LMs could memorize knowledge through text-only learning and make such predictions led to various studies exploring LMs’ internal structures, based on the assumption that their internal representations must encode some form of real-world knowledge structure. Abdou et al. (2021) and Patel and Pavlick (2022) showed that LMs’ representations of color words resemble human color perception structures. Chen et al. (2023) and Gurnee and Tegmark (2024) demonstrated that geographical relationships between countries and cities are reflected in LMs’ representational spaces. Furthermore, Hernandez et al. (2024) and Park et al. (2024b) revealed that relationships between entities and hierarchical structures of concepts are preserved in the representational space. These studies suggest that LMs’ representations capture real-world knowledge and relationships, providing important insights into what structures LMs can encode through text-only learning and where the limits of their representational capabilities lie (Bender and Koller, 2020; Bisk et al., 2020; Merrill et al., 2021; Merullo et al., 2023; Huh et al., 2024). Our research focuses on a fundamental question: whether LMs can distinguish mentions as entities. We analyzed this mention ambiguity and variability. Our findings further strengthen existing empirical evidence that LMs can acquire specific real-world knowledge structures through text-only learning.

## 8 Conclusion

We investigated how LMs encode entity mentions in their internal representations, focusing on the process of distinguishing entities in the presence of mention ambiguity and mention variability. Through experiments with five autoregressive LMs, these models achieved higher performance in entity identification compared to baselines. Our analysis revealed the degree of isomorphism between LM representations and real-world entity-centric knowledge structures, suggesting that LMs can effectively encode discrete entity relationships through text-only training. These findings deepen our understanding of how LMs internally represent and process entity information, providing new insights into the relationship between LM representation spaces and structural knowledge.

## Limitations

There are primarily three limitations in this study. First, while the structure composed of entities is inherently language-independent (“Tokyo”, “東京”, “东京” refer to the same entity), our experiments only covered English. Therefore, it remains unclear whether the structure of entity representations is truly constructed in a language-independent manner. Important future work includes investigating whether the structure of entity representations is constructed in a language-independent way, and how such representational structures relate to practical applications (such as the success rate of cross-lingual transfer, where knowledge learned in one language is transferred to another).

Second, it is unclear what defines an adequate sample size for measuring the Purity and IP. In measuring the Purity and IP, the criterion is based on whether the entity representations are mixed with “other” entity representations. Consequently, the degree of mixing is relatively determined by the prepared data. In other words, even if the LM genuinely confuses “Barack Obama” with “Michelle Obama,” it can be said that the confusion cannot be detected if either one is absent from the data. In the experiments, we used about 160,000 sentences; however, while this is a substantial dataset, we cannot conclusively determine whether this sample size is sufficient to capture all potential entity confusions that might exist in the model.

Finally, perfect separation of representations may not be optimal or desirable across all scenarios. While our study assumes that entity-level representations should be distinct, there are scenarios where representation overlap could be beneficial. For example, when treating categories like painters, musicians, and novelists as single classes, representation overlap might be advantageous. Such overlap suggests that representations cluster around broader categories like “artist,” potentially encoding hierarchical categorical knowledge. Furthermore, even in entity-level representation separation, it is debatable whether achieving perfect entity identification is a desirable property in all cases. Taking entity identification to the extreme, a model that maps all mentions of an entity to exactly the same embedding would achieve perfect clustering scores. In other contexts, the phenomenon where internal representations of instances within a class converge to a single point has been widely observed and is known as “neural collapse” (Hui et al., 2022).

Research has shown that neural collapse during training isn’t necessarily beneficial for all types of generalization and can sometimes be counterproductive. While it’s unclear whether these findings directly apply to entity representations, if LMs use classification ( $\approx$  entity identification) as an intermediate step for next-token prediction, we might expect to see neural collapse patterns emerge in intermediate layers. This raises the question of whether perfect entity identification is truly desirable. A critical direction for future work would be to investigate the potential drawbacks of having model representations that closely mirror real-world entity structures.

## Ethics Statement

In this study, we perform a novel analysis and investigate how LMs recognize knowledge about the world through internal representations. In recent years, there has been a growing concern about the risk of socially harmful biases (e.g., racial or gender biases) in the text generated by LMs. This study has the potential to contribute to a better understanding of this issue, by analyzing how LMs internally represent undesirable biases.

## Acknowledgements

We would like to thank the members of the Tohoku NLP Group for their insightful comments. This work was supported by JSPS KAKENHI Grant Number JP21K17814, JP25KJ0628, JP22H05106, JP22H03654; JST CREST Grant Number JPMJCR20D2, Japan; JST ACT-X Grant Number JPMJAX200S, Japan; and JST FOREST Grant Number JPMJFR2331, Japan.

## References

- Mostafa Abdou, Artur Kulmizev, Daniel Hershcovich, Stella Frank, Ellie Pavlick, and Anders Søgaard. 2021. [Can language models encode perceptual structure without grounding? a case study in color](#). In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 109–132, Online. Association for Computational Linguistics.
- Guillaume Alain and Yoshua Bengio. 2017. [Understanding intermediate layers using linear classifier probes](#).
- Emily M. Bender and Alexander Koller. 2020. [Climbing towards NLU: On meaning, form, and understanding in the age of data](#). In *Proceedings of the 58th Annual*

- Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. 2020. [Experience grounds language](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8718–8735, Online. Association for Computational Linguistics.
- 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.
- Xingyu Cai, Jiayi Huang, Yuchen Bian, and Kenneth Church. 2021. [Isotropy in the contextual embedding space: Clusters and manifolds](#). In *International Conference on Learning Representations*.
- Boxi Cao, Hongyu Lin, Xianpei Han, Fangchao Liu, and Le Sun. 2022. [Can prompt probe pretrained language models? understanding the invisible risks from a causal view](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5796–5808, Dublin, Ireland. Association for Computational Linguistics.
- Yida Chen, Yixian Gan, Sijia Li, Li Yao, and Xiaohan Zhao. 2023. [More than correlation: Do large language models learn causal representations of space?](#) *Preprint*, arXiv:2312.16257.
- Emily Cheng, Corentin Kervadec, and Marco Baroni. 2023. [Bridging information-theoretic and geometric compression in language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12397–12420, Singapore. Association for Computational Linguistics.
- Roi Cohen, Mor Geva, Jonathan Berant, and Amir Globerson. 2023. [Crawling the internal knowledge-base of language models](#). In *Findings of the Association for Computational Linguistics: EAACL 2023*, pages 1856–1869, Dubrovnik, Croatia. Association for Computational Linguistics.
- Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. 2022. [Knowledge neurons in pretrained transformers](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8493–8502, Dublin, Ireland. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, and 516 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, Roger Grosse, Sam McCandlish, Jared Kaplan, Dario Amodei, Martin Wattenberg, and Christopher Olah. 2022. [Toy models of superposition](#). *Preprint*, arXiv:2209.10652.
- Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. 2023. [Dissecting recall of factual associations in auto-regressive language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12216–12235, Singapore. Association for Computational Linguistics.
- Asma Ghandeharioun, Avi Caciularu, Adam Pearce, Lucas Dixon, and Mor Geva. 2024. [Patchscopes: A unifying framework for inspecting hidden representations of language models](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Wes Gurnee and Max Tegmark. 2024. [Language models represent space and time](#). In *The Twelfth International Conference on Learning Representations*.
- Lushan Han, Abhay L. Kashyap, Tim Finin, James Mayfield, and Jonathan Weese. 2013. [UMBC\\_EBIQUITY-CORE: Semantic textual similarity systems](#). In *Second Joint Conference on Lexical and Computational Semantics (\*SEM), Volume 1: Proceedings of the Main Conference and the Shared Task: Semantic Textual Similarity*, pages 44–52, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Benjamin Heinzerling and Kentaro Inui. 2021. [Language models as knowledge bases: On entity representations, storage capacity, and paraphrased queries](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1772–1791, Online. Association for Computational Linguistics.
- Benjamin Heinzerling and Kentaro Inui. 2024. [Monotonic representation of numeric attributes in language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 175–195, Bangkok, Thailand. Association for Computational Linguistics.
- Evan Hernandez, Arnab Sen Sharma, Tal Haklay, Kevin Meng, Martin Wattenberg, Jacob Andreas, Yonatan Belinkov, and David Bau. 2024. [Linearity of relation decoding in transformer language models](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.



- John Hewitt and Christopher D. Manning. 2019. [A structural probe for finding syntax in word representations](#). 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 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- Lawrence J. Hubert and Phipps Arabie. 1985. [Comparing partitions](#). *Journal of Classification*, 2:193–218.
- Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. 2024. [Position: The platonic representation hypothesis](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Like Hui, Mikhail Belkin, and Preetum Nakkiran. 2022. [Limitations of neural collapse for understanding generalization in deep learning](#). *Preprint*, arXiv:2202.08384.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.
- Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2020. [How can we know what language models know?](#) *Transactions of the Association for Computational Linguistics*, 8:423–438.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Nikolaus Kriegeskorte, Marieke Mur, and Peter Bannettini. 2008. [Representational similarity analysis - connecting the branches of systems neuroscience](#). *Frontiers in Systems Neuroscience*, 2.
- V. I. Levenshtein. 1966. Binary Codes Capable of Correcting Deletions, Insertions and Reversals. *Soviet Physics Doklady*, 10:707.
- Bastien Li  tard, Mostafa Abdou, and Anders S  gaard. 2021. [Do language models know the way to Rome?](#) In *Proceedings of the Fourth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 510–517, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Raja Marjeh, Iliia Sucholutsky, Pol van Rijn, Nori Jacoby, and Thomas L. Griffiths. 2023. [Large language models predict human sensory judgments across six modalities](#). *Preprint*, arXiv:2302.01308.
- Kevin Meng, David Bau, Alex J Andonian, and Yonatan Belinkov. 2022. [Locating and editing factual associations in GPT](#). In *Advances in Neural Information Processing Systems*.
- William Merrill, Yoav Goldberg, Roy Schwartz, and Noah A. Smith. 2021. [Provable limitations of acquiring meaning from ungrounded form: What will future language models understand?](#) *Transactions of the Association for Computational Linguistics*, 9:1047–1060.
- Jack Merullo, Louis Castricato, Carsten Eickhoff, and Ellie Pavlick. 2023. [Linearly mapping from image to text space](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Tomas Mikolov, Edouard Grave, Piotr Bojanowski, Christian Puhresch, and Armand Joulin. 2018. [Advances in pre-training distributed word representations](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013. [Linguistic regularities in continuous space word representations](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 746–751, Atlanta, Georgia. Association for Computational Linguistics.
- Marcel Milich and Alan Akbik. 2023. [ZELDA: A comprehensive benchmark for supervised entity disambiguation](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2061–2072, Dubrovnik, Croatia. Association for Computational Linguistics.
- Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter. 2020. [Zoom in: An introduction to circuits](#). *Distill*. <https://distill.pub/2020/circuits/zoom-in>.
- Core Francisco Park, Andrew Lee, Ekdeep Singh Lubana, Yongyi Yang, Maya Okawa, Kento Nishi, Martin Wattenberg, and Hidenori Tanaka. 2024a. [Icl: In-context learning of representations](#). *Preprint*, arXiv:2501.00070.
- Kiho Park, Yo Joong Choe, Yibo Jiang, and Victor Veitch. 2024b. [The geometry of categorical and hierarchical concepts in large language models](#). In *ICML 2024 Workshop on Mechanistic Interpretability*.
- Roma Patel and Ellie Pavlick. 2022. [Mapping language models to grounded conceptual spaces](#). In *International Conference on Learning Representations*.
- Fabio Petroni, Tim Rockt  schel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language models as knowledge bases?](#) 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 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Nina Poerner, Ulli Waltinger, and Hinrich Schütze. 2020. [E-BERT: Efficient-yet-effective entity embeddings for BERT](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 803–818, Online. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Jacob Mitchell Springer, Suhas Kotha, Daniel Fried, Graham Neubig, and Aditi Raghunathan. 2024. [Repetition improves language model embeddings](#). *Preprint*, arXiv:2402.15449.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Lucrezia Valeriani, Diego Doimo, Francesca Cuturello, Alessandro Laio, Alessio Ansuini, and Alberto Cazzaniga. 2023. [The geometry of hidden representations of large transformer models](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- 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.
- Ying Zhao and George Karypis. 2001. [Criterion functions for document clustering experiments and analysis](#).
- Yichu Zhou and Vivek Srikumar. 2021. [DirectProbe: Studying representations without classifiers](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5070–5083, Online. Association for Computational Linguistics.

## A Details of the Purity and Inverse Purity Algorithm

The calculation of purity (Zhao and Karypis, 2001) and inverse purity are as follows. Both metrics evaluate the quality of clusters with respect to their class distributions, with purity focusing on how well each cluster contains a single class, and inverse purity measuring how well the embeddings of each class are separated into distinct clusters.<sup>12</sup>

**Local Purity** We begin by calculating the purity for each cluster, referred to as the local purity. For each cluster  $\text{Cluster}(e)$ , we compute its local  $\text{Purity}(e)$  by following these steps:

1. For each entity  $e \in \mathcal{E}$  (e.g., "Biden"), compute the centroid  $\mathbf{b}_e \in \mathbb{B}$  of its mention embeddings  $\mathbf{X}_e = \{\mathbf{x}_e^1, \mathbf{x}_e^2, \dots\}$ :

$$\mathbf{b}_e := \frac{1}{|\mathbf{X}_e|} \sum_{\mathbf{x}_e \in \mathbf{X}_e} \mathbf{x}_e. \quad (1)$$

2. For each embedding  $\mathbf{x}$ , determine the  $\text{Cluster}(e)$  it belongs to by assigning it to the nearest centroid  $\mathbf{b}_e$ :

$$\text{Cluster}(e) := \{\mathbf{x} \mid \mathbf{b}_e = \arg \min_{\mathbf{b}' \in \mathbb{B}} d(\mathbf{x}, \mathbf{b}')\}, \quad (2)$$

where  $d(\mathbf{x}, \mathbf{y})$  represents the distance function. In our experiments, we employed Euclidean distances, which are commonly used in NLP. This step is equivalent to an intermediate stage of k-means clustering.

3. Define the gold class for entity  $e$  as:

$$\text{Class}(e) := \mathbf{X}_e. \quad (3)$$

<sup>12</sup>In Figure 3, the red-shaded area is referred to as the "class," and the blue-shaded area is referred to as the "cluster," which are used to calculate Equations 4 and 6. We calculate the local purity and local IP by considering each embedding's "class" as the entity it refers to and each embedding's "cluster" as the nearest centroid. In other words, the "clusters" refer to those formed by the Voronoi diagram using each centroid. The F1 score, composed of Purity and IP, can be used to evaluate the degree to which embeddings of entities belonging to the same class are grouped into a single cluster.

4. Compute the local purity of the  $\text{Cluster}(e)$  by calculating the fraction of embeddings in the cluster that belong to the most frequent class  $\hat{e}$ :

$$\text{local Purity}(e) := \frac{|\text{Cluster}(e) \cap \text{Class}(\hat{e})|}{|\text{Cluster}(e)|}, \quad (4)$$

where  $\text{Class}(\hat{e})$  is the most frequent class in the cluster, and  $|\text{Cluster}(e)|$  is the total number of embeddings in the  $\text{Cluster}(e)$ .

**Purity** The overall purity, which measures how well the clusters contain a single class, is calculated as the weighted average of the local purities:

$$\text{Purity}(\mathcal{E}) := \frac{1}{N} \sum_{e \in \mathcal{E}} \text{local Purity}(e) |\text{Cluster}(e)|, \quad (5)$$

where  $N$  is the total number of mention embeddings across all clusters.

**Local Inverse Purity (Local IP)** Next, we calculate the inverse purity for each entity, referred to as the local inverse purity. For each entity  $e$ , we compute local  $\text{IP}(e)$  by following these steps:

1. We use the definitions of  $\text{Cluster}(e)$  from Equation 2 and  $\text{Class}(e)$  from Equation 3 for each entity  $e$ .
2. Compute the local inverse purity local  $\text{IP}(e)$  by calculating the fraction of embeddings from the gold class  $\text{Class}(e)$  that are contained within the cluster  $\text{Cluster}(e)$ :

$$\text{local IP}(e) := \frac{|\text{Cluster}(e) \cap \text{Class}(e)|}{|\text{Class}(e)|}, \quad (6)$$

where  $|\text{Class}(e)|$  is the total number of embeddings in the gold class  $\text{Class}(e)$ .

**Inverse Purity (IP)** The overall inverse purity, representing the average degree of separation across all classes, is calculated as:

$$\text{IP}(\mathcal{E}) := \frac{1}{N} \sum_{e \in \mathcal{E}} \text{local IP}(e) |\text{Class}(e)|, \quad (7)$$

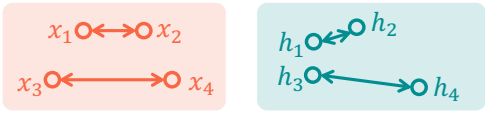

Structure	Requirement: Desired correspondence between real-world structure and representation space structure	Soft Evaluation
Metric Space (Abdou et al., 2021)	$\forall x_1, x_2, x_3, x_4 \in X$ $d_X(x_1, x_2) < d_X(x_3, x_4)$ $\Leftrightarrow d_X(h_1, h_2) < d_X(h_3, h_4)$ $\forall x_1, x_2, x_3, x_4 \in X$ $\mathbb{1}[d_X(x_1, x_2) < d_X(x_3, x_4)]$ $= \mathbb{1}[d_X(h_1, h_2) < d_X(h_3, h_4)]$	$\mathbb{P}_{x_1, x_2, x_3, x_4 \sim X} [$ $\mathbb{1}[d_X(x_1, x_2) < d_X(x_3, x_4)]$ $= \mathbb{1}[d_X(h_1, h_2) < d_X(h_3, h_4)]$ $]$ This roughly corresponds to Kendall’s $\tau$ , which in turn approximately corresponds to RSA (Kriegeskorte et al., 2008).
	 <p>Whether distance relationships are preserved</p>	
Discrete Topology (Ours)	$\forall x_1, x_2$ $x_1 = x_2$ $\Leftrightarrow h_1 = h_2$ $\forall x_1, x_2$ $\mathbb{1}[x_1 = x_2]$ $= \mathbb{1}[h_1 = h_2]$	$\mathbb{P}_{x_1, x_2 \sim X} [$ $\mathbb{1}[x_1 = x_2]$ $= \mathbb{1}[\text{Cluster}(h_1) = \text{Cluster}(h_2)]]$ This approximately corresponds to the F1 score composed of Purity and IP. By introducing clustering, we relax the requirement of exact position matching to proximity-based matching.
	 <p>Whether the same entity occupies the same position</p>	

Table 1: Comparison between existing Isomorphic analysis and our approach.

where  $N$  is the total number of mention embeddings.

This method provides a simple yet effective way to measure the embedding space in terms of both class containment (purity) and class separation (inverse purity).

**F1 Score with Purity and Inverse Purity** To further evaluate the embedding space, we calculate an F1 score using both purity and inverse purity. The F1 score provides a balanced measure that considers both the ability of the clusters to contain single classes (purity) and the degree of separation of classes into distinct clusters (inverse purity). The F1 score is computed as the harmonic mean of purity and inverse purity, formulated as follows:

$$\text{F1 Score}(\mathcal{E}) := 2 \cdot \frac{\text{Purity}(\mathcal{E}) \cdot \text{IP}(\mathcal{E})}{\text{Purity}(\mathcal{E}) + \text{IP}(\mathcal{E})}. \quad (8)$$

Here, the F1 score effectively balances the trade-off between purity and inverse purity. A high F1 score indicates that the clusters not only contain embeddings from the same class but also that embeddings from each class are well-separated into different clusters.

## B Details of the Experimental Setup

### B.1 Pre-trained Language Models

In our experiments, we used the models shown in Table 3. We used the model provided in <https://github.com/huggingface/transformers>.

### B.2 Probe Training Hyperparameters

In § 5.1, we performed linear probing using a single linear layer. The hyperparameters used for the probe training are shown in Table 4. We split the data into training and test sets with a ratio of 8:2 using stratified sampling, and performed 3-fold cross-validation with early stopping monitoring the training loss.

### B.3 Statistical Information of Experimental Data

The statistical information of the data used in the experiments is shown in Table 5.

### B.4 Detailed Definition of Mention Ambiguity

Mention ambiguity is calculated as the entropy of the frequency distribution of entity candidates for a given mention. Specifically, let  $N$  be the number

DIFFICULTY	Low-difficulty examples	High-difficulty examples
<b>Mention ambiguity:</b> Entropy of mention to entity candidate mapping.	<b>Mention:</b> ‘Ohio’  <b>Entity candidate:</b> {Ohio: 198, Ohio River: 12}	<b>Mention:</b> ‘Oxford’  <b>Entity candidate:</b> {Oxford: 26, University of Oxford: 13, Oxford, Mississippi: 6, Oxford (UK Parliament constituency): 6}
<b>Mention variability:</b> Average surface form dissimilarity of pairs of mention candidate owned by an entity.	<b>Entity:</b> Emmy Awards  <b>Mention candidate:</b> (‘Emmy’, ‘Emmy Awards’, ‘Emmy award’, ‘Emmy awards’, ‘Emmys’)	<b>Entity:</b> Māori All Blacks  <b>Mention candidate:</b> (‘New Zealand Maori Rugby’, ‘Māori All Blacks’, ‘NZ Maori’, ‘New Zealand Māori rugby union team’, ...)

Table 2: Difficulty factors considered to potentially influence entity identification, with definitions and examples. Numbers in the entity candidate lists represent the frequency of occurrence for each entity. Detailed definitions of these factors are provided in Appendix B.4, B.5.

Models	Hidden dim.	#Layer	#Head
GPT-2	768	12	12
Llama-2-7B	4096	32	32
Llama-2-13B	5120	40	40
Llama-3-8B	4096	32	32
Mistral-7B-v0.3	4096	32	32

Table 3: Hyperparameters of each model’s architecture.

Optimizer	Adam (Kingma and Ba, 2015)
Learning rate	$1.0 \times 10^{-3}$ (constant)
Batch size	1,024
Epochs	1,000

Table 4: Hyperparameters of the trained probe.

of entity candidates corresponding to a mention, and  $F_i$  be the frequency of occurrence for each candidate  $i$ . The mention ambiguity is then calculated through the following steps.

First, the total frequency  $F_{\text{total}}$  of all entity candidates for a given mention is expressed as:

$$F_{\text{total}} = \sum_{i=1}^N F_i \quad (9)$$

Next, the relative frequency  $p_i$  for each entity candidate  $i$  is defined as the ratio of the candidate’s frequency  $F_i$  to the total frequency  $F_{\text{total}}$ :

$$p_i = \frac{F_i}{F_{\text{total}}}, \quad i = 1, 2, \dots, N \quad (10)$$

Finally, mention ambiguity is represented by the entropy  $H$ , defined by the following equation:

$$H = - \sum_{i=1}^N p_i \log p_i \quad (11)$$

Using this entropy measure allows for quantitative evaluation of the degree of ambiguity associated with a given mention. When mention ambiguity is high, the difficulty of identifying the appropriate entity increases. Conversely, when mention ambiguity is low, it indicates that the corresponding entity candidates are clear and there is little ambiguity.

## B.5 Detailed Definition of Mention Variability

Mention variability is a metric that quantifies the surface-level diversity of multiple mentions referring to the same entity. In this research, we formulate this diversity as a dissimilarity measure based on Levenshtein distance. Given a set of mentions  $M_e = m_1, m_2, \dots, m_{|M_e|}$  for entity  $e$ , the surface form dissimilarity  $D$  is defined as follows:

$$D = 1 - \frac{2}{|M_e|(|M_e| - 1)} \sum_{i=1}^{|M_e|-1} \sum_{j=i+1}^{|M_e|} \left( 1 - \frac{L(m_i, m_j)}{\max(|m_i|, |m_j|)} \right) \quad (12)$$

$$= \frac{2}{|M_e|(|M_e| - 1)} \sum_{i=1}^{|M_e|-1} \sum_{j=i+1}^{|M_e|} \frac{L(m_i, m_j)}{\max(|m_i|, |m_j|)} \quad (13)$$

Here,  $L(m_i, m_j)$  represents the Levenshtein distance between mentions  $m_i$  and  $m_j$ , and  $|m_i|$  denotes the length of mention  $m_i$ . The derivation of the Levenshtein distance is described in § B.6. Using this surface form dissimilarity  $D$ , we can quantitatively evaluate the surface-level diversity (mention variability) of mentions associated with a



ZELDA-TRAIN Corpus	# Entities	# Mentions	# Sentences
All data	7,877	24,289	164,516
Ambiguous mention Subset	1,386	602	30,082

Table 5: Statistics of the experiment dataset.

given entity. A higher value of  $D$  indicates greater surface-level diversity of mentions. Conversely, a lower value of  $D$  suggests lower mention diversity and more similar surface forms.

## B.6 Derivation of the Levenshtein Distance

The Levenshtein distance (Levenshtein, 1966) measures the dissimilarity between two strings, i.e., the higher the value, the less similar the pair. This distance is calculated as the minimum number of single-character edits (comprising insertions, deletions, or substitutions) necessary to change one string into the other. Given two strings  $a$ ,  $b$  (of length  $|a|$  and  $|b|$  respectively), the Levenshtein distance  $L(a, b)$  is derived as follows:

$$L(a, b) = 1 + \min \begin{cases} L(\text{tail}(a), b) \\ L(a, \text{tail}(b)) \\ L(\text{tail}(a), \text{tail}(b)) \end{cases} \quad (14)$$

Where the  $\text{tail}(s)$  is the strings composed of the string following the except for the first character of the strings  $s$ . Note that when  $|a| = 0$ ,  $L(a, b)$  returns  $|b|$ , and when  $|b| = 0$ ,  $L(a, b)$  returns  $|a|$ . Additionally, if  $s[n]$  is the  $n$ th character of the string  $s$ , and  $a[0] = b[0]$ , the  $L(a, b)$  returns  $L(\text{tail}(a), \text{tail}(b))$ .

## B.7 Entity Representation Structural Similarity

We conduct Representational Similarity Analysis (RSA) to examine the structural similarity of entity representations across different language models and compare them against baseline embeddings. This analysis reveals how different models organize semantic knowledge about entities in their representational space.

### B.7.1 Embedding Extraction Setup

**Model Selection:** We analyze entity representations from seven different embedding sources: Llama-2-7B (layer 8), Llama-2-13B (layer 9), Llama-3-8B (layer 8), Mistral-7B (layer 8), and GPT-2 (layer 6). We use FastText wiki-news-300d-1M-subword and random embeddings as baseline comparisons.

**Layer Selection:** For transformer-based models, we extract representations from the layers that achieved the highest F1 scores in our main experiments (§4).

**Dimensionality Reduction:** All embeddings are projected to a common 20-dimensional space using Latent Dirichlet Allocation (LDA) to enable fair comparison across models with different native dimensionalities.

**Data Preprocessing:** We use the Zelda dataset. Due to computational memory constraints, we randomly sample 100,000 instances from the dataset for the RSA experiments.

### B.7.2 RSA Computation Methodology

**Representational Similarity Matrix (RSM) Construction:** For each model’s entity embeddings  $\mathbf{E} \in \mathbb{R}^{n \times d}$ , where  $n$  is the number of entity mentions and  $d = 20$  is the embedding dimension, we compute pairwise similarities based on Euclidean distances:

$$\text{RSM}_{ij} = 1 - \frac{\|\mathbf{e}_i - \mathbf{e}_j\|_2}{\max_{k,l} \|\mathbf{e}_k - \mathbf{e}_l\|_2} \quad (15)$$

**RSA Score Calculation:** We compute the representational similarity between models by correlating the upper triangular portions of their RSMs (excluding diagonal elements). We use Spearman rank correlation to capture monotonic relationships while being robust to outliers:

$$\text{RSA}(M_1, M_2) = \rho_{\text{Spearman}}(\text{flatten}(\text{triu}(\text{RSM}_1)), \text{flatten}(\text{triu}(\text{RSM}_2))) \quad (16)$$

where  $\text{triu}()$  extracts the upper triangular matrix and  $\text{flatten}()$  converts it to a vector.

## B.8 Entity Disambiguation with Patchscopes

**Model Configuration** We conduct our entity disambiguation experiments using Llama-2-13B.

**Data Preparation** Our experimental data consists of:

- **Input data:** Following the main experiments in § 4, we use approximately 7,800 entities and 160,000 mention embeddings.

- **Embedding configuration:** 20-dimensional LDA embeddings.

**Patchscopes Configuration** We employ the Patchscopes (Ghandeharioun et al., 2024) with the following specifications:

- **Source representation:** Final token representation from the entity mention in context, consistent with our main experiments in § 4.
- **Target prompt:** Binary choice format: 10-shot examples + “Does X refer to A or B? One-word answer only. A: {option\_a} B: {option\_b} Answer:”
- **Intervention layers:** We analyze all layers to understand how entity representations evolve across model depth.
- **Position indices:**
  - Source position: -1 (last token of the mention)
  - Target position: Automatically detected X token in the prompt template

**Prompt Design** We use a standardized prompt template with 10-shot in-context learning examples. This prompt design aims to constrain the model’s output to either “A” or “B”. By showing the model 10 examples, we successfully guide it to output only “A” or “B” as answers.

Does Michael Jordan refer to A or B?  
One-word answer only.  
A: Spain  
B: Michael Jordan  
Answer: B  
[... 9 more examples ...]  
Does X refer to A or B?  
One-word answer only.  
A: {option\_a}  
B: {option\_b}  
Answer:

The assignment of correct entity vs. distractor to options A/B is deterministically randomized using MD5 hashing of the mention-entity pair to ensure balanced label distribution.

**Distractor Selection** We select distractors as the nearest neighbor entities to X’s embedding (contextualized mention embeddings from Llama-2-13B). An actual example is shown below:

Does X refer to A or B?  
One-word answer only.  
A: San Diego  
B: San Francisco  
Answer:

This approach is based on the intuition that if the representations are mixed with those of neighboring entities, the entity disambiguation task will likely fail. Using nearest neighbors as distractors provides categorically similar alternatives, creating an appropriately challenging task compared to random distractor selection. For instance, the nearest neighbor to “Tesla, Inc.” is “Nissan”, both being automobile manufacturers, making the discrimination task non-trivial yet meaningful.

**Baselines** The baselines are chance rate and w/out PatchScope. Since this task is binary choice, the chance rate is 50%. w/out PatchScope does not patch entity representations to “X”. That is, it inputs the example and Does X refer to A or B? One-word answer only. A: San Diego B: San Francisco Answer: for in-context learning (ICL). w/out PatchScope is expected to perform nearly at chance level, and indeed, Figure 8 (a) confirms that it performs equivalently to the chance rate.

## C Details of the Analysis Results

### C.1 Dimensions and F1 Score

The correspondence between the number of dimensions and F1 scores for each model is shown in Table 6.

### C.2 Mention Ambiguity

#### C.2.1 Model Comparison Results

Table 7 shows the AUC of mention ambiguity and F1 score for the baseline. Table 8 and Figure 12 show the analysis results from layer 0 to the final layer for mention ambiguity using GPT-2, Llama-2 7B, Llama-2 13B, Llama-3 8B, and Mistral 7B.

#### C.2.2 Results with Average Subword Token Embeddings

Figure 9 shows the results when using average embeddings of subword tokens as the target tokens for analysis. We can observe that the trends are nearly identical to those in Figure 5 in § 4, which used last token embeddings.

Dimension	Llama-2-7B	Llama-2-13B	Llama-3-8B	Mistral-7B	GPT-2	FastText	Random Embeddings
1	0.01	0.01	0.00	0.01	0.00	0.23	0.00
2	0.07	0.07	0.06	0.06	0.05	0.38	0.00
5	0.51	0.51	0.52	0.49	0.46	0.68	0.00
10	0.80	0.80	0.81	0.78	0.77	0.80	0.00
20	0.91	0.91	0.91	0.90	0.85	0.85	0.01
30	0.93	0.93	0.93	0.92	0.87	0.87	0.01
40	0.94	0.94	0.94	0.93	0.87	0.87	0.02
50	0.94	0.94	0.94	0.94	0.88	0.88	0.05
60	0.95	0.95	0.95	0.94	0.88	0.88	0.05
80	0.95	0.95	0.95	0.95	0.88	0.88	0.08
100	0.95	0.95	0.95	0.95	0.88	0.88	0.10
200	0.96	0.96	0.95	0.96	0.88	0.88	0.30
300	0.96	0.96	0.95	0.96	0.88	0.85	0.48
500	0.96	0.96	0.96	0.96	0.90	–	0.69
1000	0.96	0.96	0.96	0.96	–	–	0.87
2000	0.96	0.96	0.96	0.96	–	–	0.96
3000	0.96	0.96	0.96	0.96	–	–	0.98
4096	0.93	0.93	0.93	0.94	–	–	0.99

Table 6: F1 scores corresponding to the number of dimensions for each model.

Method	Mention Ambiguity	Mention Variability
Random Embeddings	0.666	0.003
Unique Mention Embeddings	0.375	0.413
FastText	0.352	0.611

Table 7: AUC scores for baseline embedding methods. These static embeddings serve as comparison points for the contextualized representations from transformer models shown in Tables 8 and 9. All embeddings are reduced to 20 dimensions.

### C.3 Mention Variability

#### C.3.1 Model Comparison Results

Table 7 shows the AUC of mention variability and F1 score for the baseline. Table 9 and Figure 13 show the analysis results from layer 0 to the final layer for mention variability using GPT-2, Llama-2 7B, Llama-2 13B, Llama-3 8B, and Mistral 7B.

#### C.3.2 Results with Average Subword Token Embeddings

Figure 9 shows the results when using average embeddings of subword tokens as the target tokens for analysis. The cross-model trends remain consistent with those observed in Figure 5 (§ 4). Notably, however, Llama-2 7B at layer 0 (bottom right in Figure 9) achieves higher F1 scores compared to the last token embedding results in Figure 5. This improvement likely stems from the averaging process enabling greater reliance on surface-level character information for entity identification.

### C.4 Qualitative Analysis of Output Inconsistency

We conducted a qualitative analysis to investigate the impact of mention variability on model outputs. Interestingly, even when outputs were inconsistent across different surface forms, the predicted years in patterns like “[X] was founded in [Y]” remained temporally proximate. For example:

- “Red Crescent Society was founded in 1934” vs. “Red Crescent was founded in 1919” (15-year difference)
- “Holy Cross Brothers was founded in 1847” vs. “Brothers of Holy Cross was founded in 1845” (2-year difference)

This suggests that while models struggle with exact consistency across mention variants, they maintain approximate temporal coherence, capturing the general historical period rather than memorizing exact dates.

### D Supplementary Experiment: Adjusted Rand Index for Entity Identification

In addition to the Purity and Inverse Purity (IP) metrics presented in the main body of this paper, we conducted a supplementary experiment using the Adjusted Rand Index (ARI) (Hubert and Arabie, 1985) to evaluate the extent to which our models’ internal representations facilitate entity identification. The Adjusted Rand Index is a measure of the similarity between two data clusterings, accounting

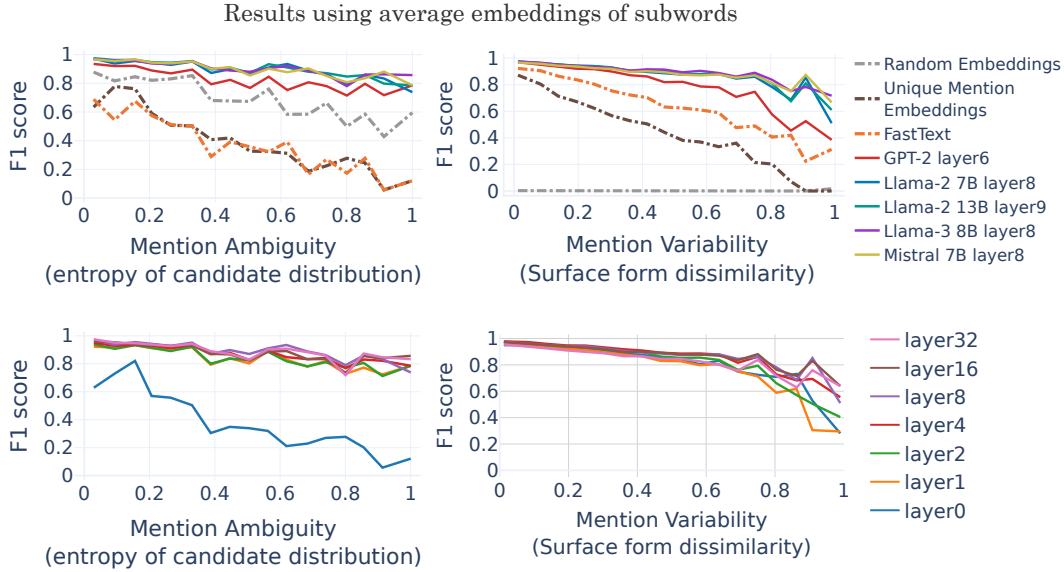


Figure 9: Separability (F1 scores) of entity representations using average embeddings of subword tokens for each difficulty (left: mention ambiguity, right: mention variability). F1 scores are averaged within each bin on the x-axis. Top: Results across all models, using the layer with highest AUC for LM scores. Bottom: Layer-wise analysis for Llama-2 7B.

for chance. In this context, it serves as an alternative approach to assessing the quality of entity identification, complementing our primary metrics.

### D.1 Experimental Setup

The experiment was designed to quantify how well mentions of the same entity cluster together in the embedding space. We used the same dataset and entity representation definitions as described in Section 3. Cluster formation followed the same methodology as detailed in Section 2.2, primarily forming clusters based on the entity of the nearest centroid to a given mention embedding. The Adjusted Rand Index was then computed between these generated clusters and the gold-standard entity labels. The Adjusted Rand Index (ARI) is calculated as:

$$\text{ARI} = \frac{\text{RI} - E[\text{RI}]}{\max(\text{RI}) - E[\text{RI}]} \quad (17)$$

where RI is the Rand Index,  $E[\text{RI}]$  is the expected Rand Index under random partitioning, and  $\max(\text{RI})$  is the maximum possible Rand Index (which is 1). The Rand Index (RI) itself measures the agreement between two clusterings by considering all pairs of samples and counting pairs that are assigned in the same or different clusters in both the predicted and true clusterings. ARI corrects for chance agreement, ensuring that random labelings

have an ARI close to 0.0, and perfect agreement results in an ARI of 1.0.

### D.2 Results

The results of this supplementary experiment are presented in Figure 11. Consistent with the trends observed using Purity and IP, the Adjusted Rand Index scores reveal a largely similar relative ranking among the models. Specifically, models that performed well on Purity and IP generally also achieved higher Adjusted Rand Index scores. This consistency further validates our findings regarding the models’ ability to distinguish between different entities based on their contextual embeddings.

### E Computational Resources

For this experiment, we utilized four NVIDIA RTX 6000 Ada graphics cards.

### F Usage of AI assistants

We used an AI Assistant (Claude) for writing this paper and developing source code for the experiments. However, its use was limited to code completion, translation, text refinement, and table creation, while the content and ideas are solely those of the authors.






Figure 10: UMAP visualization of the ambiguous mention “Georgia.” In layer 0 of Llama-2 7B, the embeddings are mixed, but in layer 8, the embeddings are grouped by entity.

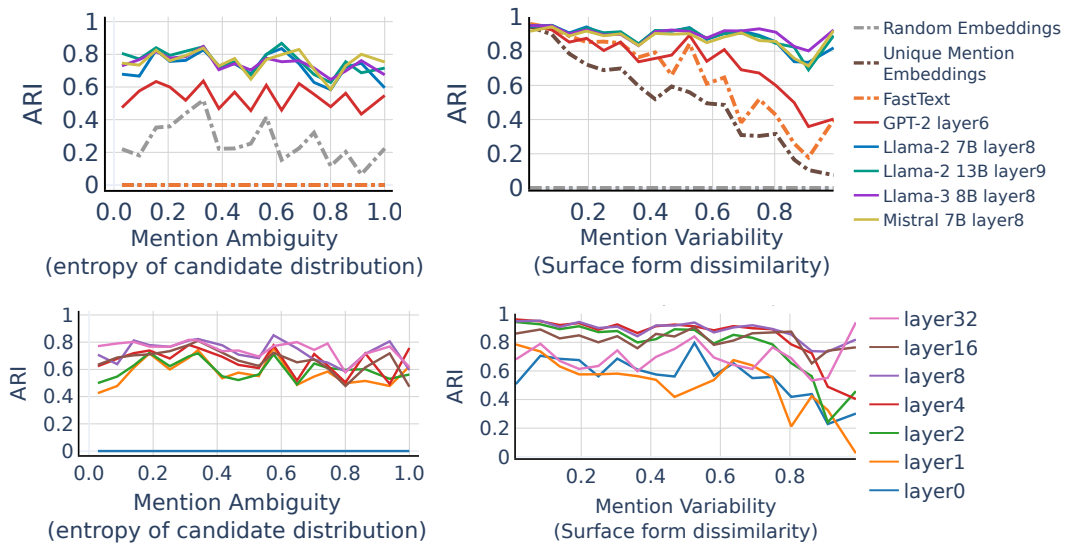


Figure 11: ARI (Adjusted Rand Index) of entity representations for each difficulty (left: mention ambiguity, right: mention variability). ARI scores are averaged within each bin on the x-axis. Top: Results across all models, using the layer with highest AUC for LM scores. Bottom: Layer-wise analysis for Llama-2 7B.

Mention Ambiguity (dim=20)					
Layer	GPT-2	Llama-2 7B	Llama-2 13B	Llama-3 8B	Mistral 7B
0	0.535	0.381	0.385	0.384	0.381
1	0.767	0.808	0.819	0.828	0.826
2	0.769	0.817	0.816	0.836	0.821
3	0.793	0.819	0.843	0.844	0.832
4	0.792	0.849	0.864	0.869	0.836
5	0.796	0.848	0.870	0.870	0.841
6	0.797	0.855	0.873	0.875	0.873
7	0.798	0.859	0.867	0.875	0.867
8	0.793	0.868	0.881	0.878	0.879
9	0.795	0.869	0.882	0.874	0.883
10	0.795	0.868	0.888	0.876	0.884
11	0.795	0.870	0.887	0.873	0.880
12	0.801	0.866	0.881	0.876	0.880
13	-	0.867	0.870	0.869	0.886
14	-	0.854	0.869	0.866	0.876
15	-	0.850	0.871	0.866	0.869
16	-	0.844	0.873	0.851	0.862
17	-	0.841	0.865	0.850	0.857
18	-	0.844	0.865	0.848	0.851
19	-	0.846	0.865	0.848	0.848
20	-	0.844	0.859	0.840	0.851
21	-	0.850	0.859	0.839	0.850
22	-	0.852	0.867	0.846	0.854
23	-	0.850	0.871	0.841	0.855
24	-	0.860	0.870	0.839	0.855
25	-	0.857	0.872	0.840	0.853
26	-	0.862	0.874	0.840	0.858
27	-	0.866	0.875	0.846	0.863
28	-	0.871	0.874	0.849	0.860
29	-	0.868	0.874	0.850	0.861
30	-	0.871	0.872	0.855	0.868
31	-	0.873	0.871	0.860	0.866
32	-	0.877	0.871	0.878	0.877
33	-	-	0.869	-	-
34	-	-	0.876	-	-
35	-	-	0.876	-	-
36	-	-	0.877	-	-
37	-	-	0.881	-	-
38	-	-	0.881	-	-
39	-	-	0.881	-	-
40	-	-	0.877	-	-

Table 8: AUC scores measuring the area under the curve where the x-axis represents mention ambiguity and the y-axis represents F1 scores derived from Purity and IP metrics across different model layers. All embeddings are reduced to 20 dimensions.

Mention Variability (dim=20)					
Layer	GPT-2	Llama-2 7B	Llama-2 13B	Llama-3 8B	Mistral 7B
0	0.377	0.290	0.290	0.441	0.305
1	0.574	0.473	0.506	0.737	0.693
2	0.589	0.695	0.748	0.805	0.780
3	0.630	0.779	0.806	0.817	0.803
4	0.644	0.795	0.825	0.822	0.826
5	0.654	0.812	0.825	0.833	0.819
6	0.663	0.815	0.839	0.825	0.826
7	0.650	0.818	0.832	0.824	0.821
8	0.658	0.817	0.831	0.831	0.816
9	0.648	0.797	0.822	0.823	0.810
10	0.646	0.794	0.810	0.824	0.802
11	0.648	0.797	0.803	0.824	0.804
12	0.645	0.791	0.794	0.822	0.802
13	-	0.784	0.782	0.811	0.795
14	-	0.781	0.771	0.811	0.788
15	-	0.766	0.764	0.799	0.783
16	-	0.761	0.773	0.787	0.764
17	-	0.748	0.758	0.776	0.754
18	-	0.729	0.753	0.768	0.751
19	-	0.723	0.743	0.756	0.737
20	-	0.711	0.722	0.742	0.725
21	-	0.701	0.721	0.725	0.715
22	-	0.692	0.710	0.730	0.704
23	-	0.693	0.701	0.718	0.692
24	-	0.686	0.693	0.717	0.690
25	-	0.682	0.691	0.723	0.689
26	-	0.673	0.688	0.722	0.680
27	-	0.670	0.681	0.727	0.678
28	-	0.665	0.680	0.724	0.682
29	-	0.657	0.675	0.726	0.681
30	-	0.649	0.674	0.737	0.685
31	-	0.647	0.672	0.738	0.678
32	-	0.635	0.668	0.665	0.662
33	-	-	0.663	-	-
34	-	-	0.657	-	-
35	-	-	0.655	-	-
36	-	-	0.651	-	-
37	-	-	0.648	-	-
38	-	-	0.653	-	-
39	-	-	0.660	-	-
40	-	-	0.655	-	-

Table 9: AUC scores measuring the area under the curve where the x-axis represents mention variability and the y-axis represents F1 scores derived from Purity and IP metrics across different model layers. All embeddings are reduced to 20 dimensions.

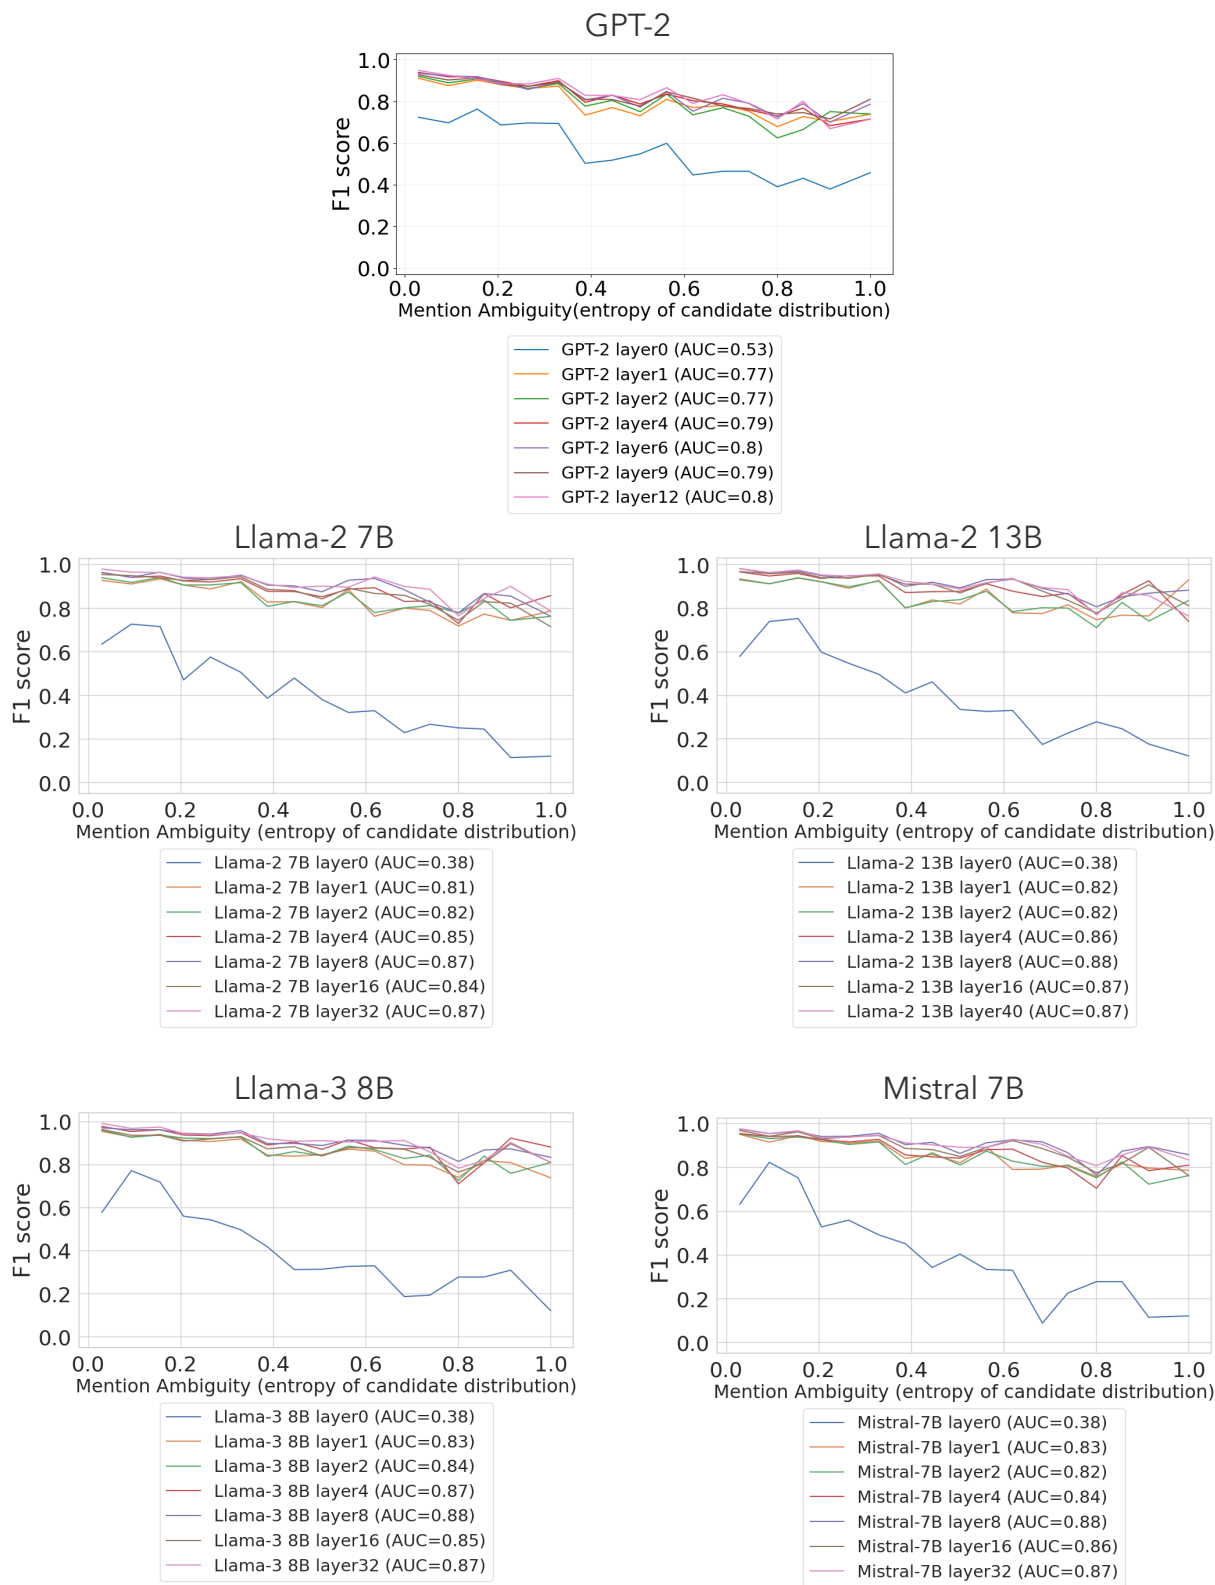


Figure 12: Results for all models: mention ambiguity and the separability of entity representations (F1 score).



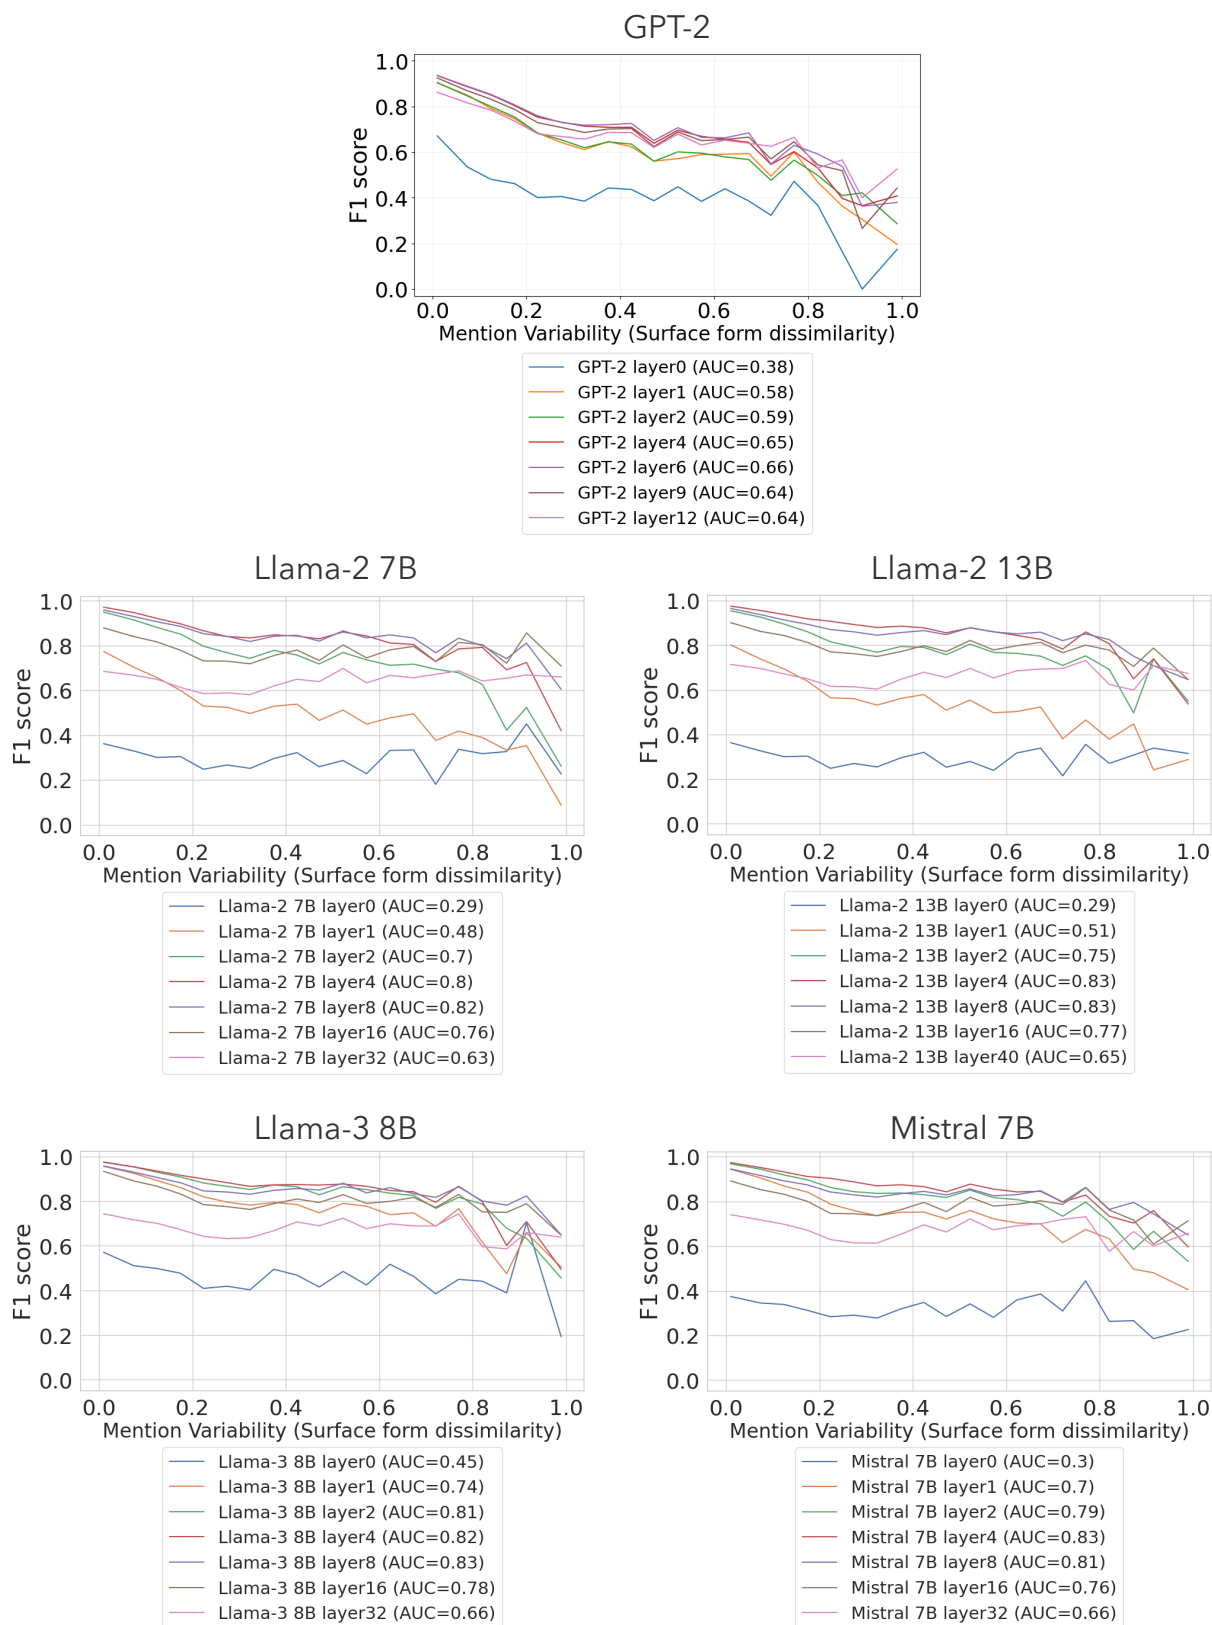


Figure 13: Results for all models: mention variability and the separability of entity representations (F1 score)