A QUANTITATIVE ANALYSIS OF SEMANTIC INFORMATION IN DEEP REPRESENTATIONS OF TEXT AND IMAGES

Anonymous authorsPaper under double-blind review

000

001

003

010 011

012

013

014

015

016

017

018

019

021

023

025

030

033

034

035

036

037

040

041

042

043

044

045

046

047

048

051

052

ABSTRACT

Deep neural networks are known to develop similar representations for semantically related data, even when they belong to different domains, such as an image and its description, or the same text in different languages. We present a method for quantitatively investigating this phenomenon by measuring the relative information content of the representations of semantically related data and probing how it is encoded into multiple tokens of large language models (LLMs) and vision transformers. Looking first at how LLMs process pairs of translated sentences, we identify inner "semantic" layers containing the most language-transferable information. We find moreover that, on these layers, a larger LLM (DeepSeek-V3) extracts significantly more general information than a smaller one (Llama3.1-8B). Semantic information of English text is spread across many tokens and it is characterized by long-distance correlations between tokens and by a causal leftto-right (i.e., past-future) asymmetry. We also identify layers encoding semantic information within visual transformers. We show that caption representations in the semantic layers of LLMs predict visual representations of the corresponding images. We observe significant and model-dependent information asymmetries between image and text representations.

1 Introduction

Current research explores the conjecture that Large Language Models (LLMs) possess a *universal language of thought*, enabling them to process concepts across different languages Peng, Søgaard (2024); Lindsey et al. (2025); Huh et al. (2024); Singh et al. (2019); Brinkmann et al. (2025). A compelling example comes from Anthropic's study of the Claude model Lindsey et al. (2025), which shows that the network occasionally reasons within a shared conceptual space across languages. Intriguingly, the extent of this cross-lingual alignment appears to grow with model size, suggesting that universality may be an emergent property of scale. Ref. Huh et al. (2024) further argues more generally that representations in deep networks, including but not limited to LLMs, are converging, potentially pointing toward shared underlying conceptual structures.

In this work, we seek to turn the qualitative notion that deep networks develop semantically meaningful shared representations into quantitative measures, identifying where in the network, and to what extent, semantic information is encoded. We can cast the identification of shared semantic content as a problem of relative information: different networks learn distributions defined on different data domains, such as images and captions. Those distributions loosely span a common underlying subspace, associated with shared semantic information. The support of the different distributions does not reduce to this subspace, but spans other subspaces, which are data and architecture specific. To quantify relative information we need a similarity criterion that is both (i) asymmetric, as there may be a partial order relation between models and representations (qualitatively, the subspaces which are data specific might have different dimensions); (ii) computationally efficient for representations of dimension of order $10^5 - 10^6$, which is the typical number of neurons in a deep representation. Cross-entropy measures how difficult it is to encode an event from a distribution into another and would be the ideal choice, but estimating it is computationally difficult due to the dimension of the representations. On the other hand, local methods like the average number of nearest neighbors Doimo et al. (2020); Huh et al. (2024) fail to capture the asymmetry between spaces. To overcome the limitations, we employ the information imbalance Glielmo et al. (2022): a method that leverages

conditional ranking to provide an asymmetric measure of relative mutual predictivity, which has been shown to be an excellent proxy of the cross entropy Tatto et al. (2024).

We analyze the representation of translations of the same sentences in different languages, using one of the most powerful open-source language models, DeepSeek-V3 with 671B parameters DeepSeek-AI et al. (2025), comparing it to the mid-sized Llama3.1-8b LLM Meta (2024). Moreover, we analyze the deep representations of pairs of images and of human-generated descriptions, and of pairs of images depicting the same object, with the goal of capturing the alignment between image and caption representations in different architectures and the relative predictive power of images and texts. As visual models, we employ image-gpt-large Chen et al. (2020) and DinoV2-large Oquab et al. (2024).

Our main contributions are as follows:

054

056

057

060

061

062

063

064 065

066

067

068

069

070

071

073

075

076

077

079

081

083

084

085

087

091

092

094 095

096

097

098

100

101

102

103 104

105

107

- We analyze the DeepSeek-V3 language model and measure, layer by layer, the relative
 information content between representations of sentence pairs that are translations of each
 other. This reveals a broad region of the network, robust across language pairs, that
 consistently encodes shared semantic content, and whose extent increases with model size.
- We show that in deep layers, information is distributed across tokens through long-range correlations. Moreover, we observe that better semantic representations present stronger token-token correlations. In the most semantically informative layers of DeepSeek-V3 processing English text, past tokens predict future tokens more accurately than the reverse, highlighting a causal asymmetry in semantic representations.
- We also identify semantic layers in vision transformers by processing pairs of images that share the same class. These semantic layers are also the most informative about DeepSeek-V3 representations of descriptive image captions, where we observe significant information asymmetries between image and text representations.

Related work. The presence of semantic layers containing shared information in deep networks is connected to the notion of representation alignment Sucholutsky et al. (2024). Two different models can be 'stitched' together with Bansal et al. (2021) or without Moschella et al. (2023) a trainable stitching layer at several different stages of the network. For instance, two vision models can effectively use similar low-level features Lenc, Vedaldi (2019). Nonetheless, the types of features that we probe for in a semantic layer are instead abstract, encoding the collective relationships between the input constituents, independently of the particular way in which they were expressed in the specific input modality at disposal. The general representations we are probing lie behind the so-called *Platonic representation hypothesis*, postulating that, as model quality improves, models converge to similar representations Huh et al. (2024). We significantly extend our understanding of the nature of these shared representations with respect to Huh et al. (2024), since for autoregressive language models, the authors take the average across the token axis, reducing the dimensionality of the problem, whereas for vision models, they only consider the "class" token. We address instead where semantic information is encoded within the network, namely across different tokens and different layers. The representation of language has also been studied in Ref. Cheng et al. (2025) using, among other techniques, the Information Imbalance. This work is focused on the analysis of the last-token representation of the same text in different models, and not on multimodal data.

Other earlier works point to deep similarities across architectures and modalities Kornblith et al. (2019); Sorscher et al. (2022); Maniparambil et al. (2024). Again, these works do not focus on determining the location and nature of the shared semantic layers. Moreover, they do not include a latest-generation, huge model such as DeepSeek-V3. Our work also relates to studies that probe the extent to which multilingual LLMs produce shared representations across languages (and whether this is due to English acting as a pivot language), e.g.: Peng, Søgaard (2024); Singh et al. (2019); Brinkmann et al. (2025); Wendler et al. (2024); Zhao et al. (2024). To this research line, we contribute a method to compute similarity that allows us to take long spans of tokens into account and to capture asymmetry between language-specific representations.

2 THE INFORMATION IMBALANCE

Given two feature spaces A and B pertaining to a collection of data points, the II is a measure that quantifies how much information on feature space B is carried by feature space A and vice-versa.

For an introduction to the II with simple examples, we refer the reader to Ref. Glielmo et al. (2022). In this work, the data points are the high-dimensional representations generated by artificial neural networks. The central idea is to compute the distance between every data point in each feature space separately, obtaining for each point i, two sets of distances $d_{i,j}^A$ and $d_{i,j}^B$, $i,j=1,...,N_s$, respectively, where N_s is the number of samples. Then, for each data point i, all other points j are ranked in each feature space in increasing order, obtaining two sets of integer ranks $r_{i,j}^A$ and $r_{i,j}^B$. The II between feature spaces A and B, denoted by $\Delta(A \to B)$, is proportional to the average rank of points in space B of the nearest neighbors in space A:

$$\Delta(A \to B) = \frac{2}{N_s - 1} \langle r^B | r^A = 0 \rangle \approx \frac{2}{N_s - 1} \frac{1}{N} \sum_{i,j} \delta_{r_{ij}^A, 0} r_{ij}^B, \tag{1}$$

where in the last step we estimate the expected value over the dataset, $\langle . \rangle$, as the sample mean, and $\delta_{r_{ij}^A,0}$ is a Kronecker delta fixing data samples i,j to be first neighbours in space A. Intuitively, if A predicts B to a large degree, the rank in space B of the first neighbors in space A should be small (equal to zero in the case of perfect predictability). Importantly, the II can be asymmetric: if for example A contains more information about B than B about A, $\Delta(A \to B) < \Delta(B \to A)$. The normalization is chosen such that if feature space A is non-informative about B $\Delta(A \to B) = 1$. While feasible to compute numerically, the II is formally connected to the theory of copula variables Glielmo et al. (2022), given that distance ranks are their discrete version, and it provides an upper bound to a conditional mutual information Glielmo et al. (2022); Tatto et al. (2024). Implementation details are in Supp. Inf., Sec. C.

In this work, we study how the information is spread across the tokens using models where the embedding dimension is of the order of a few thousands and the number of tokens ranges from 40 to 1024, bringing to high-dimensional feature spaces of dimension up to $\mathcal{O}(10^6)$. Thus, to save memory, we binarize the activations through the application of the sign function, a 'quantization method' which has been successfully applied in neural network training Hubara et al. (2018); Guo (2018); Hubara et al. (2016); Wang et al. (2023), and in the analysis of representations. For example, it has been shown that it approximately preserves the angles between weights and activations Anderson, Berg (2018), and recently it was shown to be suitable for calculating the intrinsic dimension of many-token data representations Acevedo et al. (2025). It is common practice to clip activations at a fixed quantile and normalize vectors before computing distances, in order to mitigate the effects of outlier dimensions Kovaleva et al. (2021); Bondarenko et al. (2023) and extreme activation values Sun et al. (2024). Binarization implicitly achieves both: it clips all activations to ± 1 and ensures that vectors of the same length have equal norm. For a case in which we can compute the II in full precision, we show in Sec. H of the Supp. Inf. that binarizing and computing distances with the Hamming distance leads to remarkably similar results to those obtained with normalized real-valued activations and L_2 distances.

3 RESULTS

3.1 THE LOCAL GEOMETRY OF LLM REPRESENTATIONS OF TRANSLATED SENTENCES ALLOWS THE IDENTIFICATION OF SEMANTIC LAYERS

We consider sets of features A_l and B_l in two different languages. The index l labels different layers and ranges from 1 to L. If a sentence in the language associated with space A_l is tokenized into T tokens, then at each layer l, $A_l \in \mathbb{R}^{T \times E}$, where E denotes the embedding dimension of the language model processing the sentence. What we refer to as the semantic content (or meaning) of a sentence, is a global property, emerging from the interaction of all the symbols present in it. However, in the deep representations the attention mechanism can move information across the tokens, and eventually concentrate it in a specific token, as indeed happens by construction in the last layer of causal models. One of the goals of this work is studying explicitly where the information is stored in the deep layers. This will be addressed in Section 3.1.2. We use human-made translations from Helsinki-NLP/opus_books. Data processing details are in Supp. Inf. D.

Starting with English-Spanish translations, we compute the information imbalance between A_l and B_l for every layer in the network, excluding the embedding layer, using the last 20 tokens of each representation (in Sec. 3.1.2 we study the dependence of the II on the number of tokens). Note that

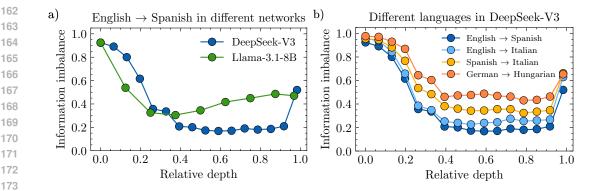


Figure 1: **Panel a)** Information Imbalance (II) from English to Spanish, using representations generated at equal depth from translated sentences of opus_books, as a function of the relative network depth, for DeepSeek-V3 and Llama3.1-8b. A smaller value of the II correspond to higher predictive power. We used the concatenation of the last 20 tokens for the computation. **Panel b)** II between equal-layer representations of DeepSeek-V3 processing different translation pairs. For each language pair in the legend, we show the II from the first language to the second. Error bars computed with 5 repetitions of a Jackknife procedure, subsampling 2000 sentences out of 5000, are smaller than marker size.

 A_l and B_l are generated independently, with the network reading one language at a time, without any specific prompt or instruction. Results are shown in panel a) of Fig. 1. Since the initial layers of LLMs must process the input, the corresponding representations are strongly language-dependent, and hence not highly informative about each other. As we consider deeper layers, we observe that they become progressively more informative about the deep representations in the other language. Since the information shared by both input languages must be semantic in nature, this analysis confirms that the data structure of deep representations seems to encode information in a universal, language-independent manner. For the very last layers of DeepSeek-V3, we observe a sudden rise in Information Imbalance (II), since those layers are generating the output in a specific language, which cannot be universal.

Note that for an ideal dataset of translations, an ideal network separately processing the sentence pairs should generate equivalent representations, namely somewhere in the network the II from one language to the other should be close to zero. Thus, better semantic representations in our experiments have a higher predictability and a lower II. In Fig. 1, for English-Spanish pairs the II of DeepSeek-V3's representations has a very broad minimum of order 0.2, roughly between relative depths 0.4 and 0.9. This number is impressively small, considering that the dimension of the representations which are compared is $20 \cdot 7000 \sim 150,000$. As a comparison, if the data were generated by Gaussian processes, one would observe an II of 0.2 if approximately 20% of the features were shared between the two representations (see Supp. Inf., Sec. C). Instead, for Llama3.1-8b the minimum II is significantly higher, of order 0.3, corresponding to approximately 15% shared features. Moreover, the minimum is narrower, around a relative depth of 0.3, suggesting that bigger and better models generate representations of higher semantic content.

We also measured the intrinsic dimensionality of DeepSeek-V3's representations using the BID estimator Acevedo et al. (2025). We find two local maxima of roughly the same height that coincide with the beginning (around 0.4) and the end (around 0.9) of the semantic region observed in Fig. 1, see Supp. Inf., Sec. I. Although these peaks suggest that semantic encoding leverages a higher-dimensional abstract space—echoing findings from brain-similarity studies Antonello, Cheng (2024)—we defer a deeper investigation of representation dimensionality to future work.

As a sanity check, we show in Supp. Inf. F that performing batch-shuffling on any of the datasets leads to completely uninformative representations (II gives 1 for every layer), since the semantic correspondence between translations is destroyed. Furthermore, in Supp. Inf. G we show that our results are compatible with those obtained using the overlap neighborhood metric from Ref. Huh et al. (2024). In Supp. Inf. E we also show that, if instead of considering many-token representations as

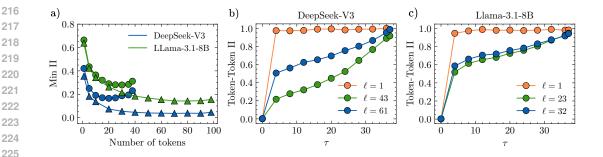


Figure 2: **Panel a**) Minimum Information Imbalance (II) across depth, from English to Spanish representations, as a function of the number of tokens used in the computation, for DeepSeek-V3 and Llama3.1-8b, when considering shorter (40 to 80 tokens, in circles) or longer (100 to 200 tokens, in triangles) sentences. **Panels b**) **and c**) II from the last token to a previous token at token-distance τ , using English sentences, computed on the representations generated by DeepSeek-V3 in panel a) and by Llama3.1-8b in panel b), in different layers. Long-distance token-token correlations are maximal (II increases most slowly) for representations on the inner semantic layers (43 for DeepSeek-V3, 23 for Llama3.1-8b), and the effect is dramatically stronger in DeepSeek-V3. Error bars computed with 5 repetitions of a Jackknife procedure, subsampling 2000 sentences out of 5000, are smaller than marker size.

we do in Fig. 1, one reduces the dimensionality by averaging all tokens, as done in Ref. Huh et al. (2024), the results are qualitatively compatible, but present significant quantitative differences. For all layers, but most strongly on the first and the last layers, we observe that averaging significantly increases the level of mutual predictability.

3.1.1 HETEROGENEITY OF DEEPSEEK'S REPRESENTATIONS FOR DIFFERENT LANGUAGE PAIRS

We repeated the previous translation experiment with other language pairs, in particular English-Italian, English-Spanish, Italian-Spanish and German-Hungarian. We considered again sentences taken from opus_books Tiedemann (2012) and applied the same filtering procedure regarding their number of tokens, as in Fig. 1a).

Fig. 1b) shows the same qualitative behavior observed for English-Spanish pairs (low information imbalance in the deep layers) but with major quantitative differences. Even if the training dataset of DeepSeek-V3 contains several languages in which the network can perform well, the quality of the representations is heterogeneous across languages, since it depends on the amount of data seen in each language. Fig. 1b) shows that the information imbalance between English and Italian is higher than between English and Spanish. We conjecture that this difference is due to the greater online presence of the Spanish language relative to Italian Lan (2025), and the consequent difference in the amount of training data. Aligned with this argument, we observe that, even though Spanish and Italian are linguistically closer than either is to English, the Information Imbalance between Spanish and Italian is higher than between English-Spanish or English-Italian translation pairs. Finally, we observe that translations involving a less common language, namely Hungarian, are associated with a significantly higher II (the minimum value is 0.5).

3.1.2 SEMANTIC INFORMATION IN LLMs is SPREAD ACROSS MANY TOKENS

Having shown that long-span Information Imbalance captures the semantic similarity of sentences in different languages, we turn our attention to the extent to which different tokens are contributing to cross-linguistic similarity. Fig. 2a) shows the minimum II across all layers as a function of the number of tokens involved in the calculation, again for English-Spanish translations from opus_books. The span is determined starting from the end and moving back towards the sentence beginning. The main result here, for both models, is that the addition of more tokens significantly increases the mutual predictivity of the translation representations, proving that semantic information is not concentrated in the last tokens, but spread over many of them.

271

272

273

274

275

276

277

278

279

280

281

282

283

284

285

286287288

289 290

291

292

293

294

295

296

297

298

299

300

301

302

303

304

305

306

307

308

309

310 311

312

313 314

315

316

317

318

319

320

321

322

323

Focusing first on DeepSeek-V3, we note that, up to about 10 tokens, the results do not change much between shorter (40 to 80 tokens) and longer (100 to 200 tokens) sentences. In both cases adding more tokens significantly boosts the information content. However, for shorter sentences we observe that the II reaches its lowest value at around 20 tokens, and it increases with the addition of further tokens. We conjecture that this is because the last added tokens are close to the start of the sentence, and thus they do not carry much information about the semantic content of the whole sentence, given the autoregressive nature of the attention mechanism in the considered language models. In contrast, for longer sentences, we observe that adding tokens leads to a rapid decrease in II, which converges to very low values (around 0.05) at around 50 tokens and stays there. On the one hand, this reflects the fact that, as longer sentences contain more information, their translation will contain more shared information at the semantic level. On the other hand, it further confirms that information is spread across a large number of tokens. We observe similar trends for Llama3.1-8b II, although the curves are systematically above the DeepSeek-V3 ones. Similar to what we observed in Fig 1a), this confirms that bigger and better performing models generate representations of better quality, capturing deep semantic similarity across different languages. Finally, this result shows that not every token contributes homogeneously to the semantic similarity between representations, and thus this analysis suggests to take the concatenation of roughly the last 20 tokens as representative of the sentence content.

3.1.3 Inner representations have longer-range correlations

The previous analysis implies that the representation of a sentence in the deep layers, where it is transferrable between different languages, is encoded in several tokens. One wonders if the activations of the tokens in those layers are uncorrelated or not. If activations were independent, one could argue that different tokens are associated to different "directions" in semantic space. If they are correlated, this would suggest that the model is building a joint semantic representation of sentences spreading its information across its component tokens. We address this question by measuring how informative two distinct tokens of the same sentence are about each other. Fig. 2 (panels b) and c)) shows how much the II between two token positions depends on their distance, for sentences in English. Concretely, we measure $\Delta(T \to T - \tau)$, the II between the token at position T, and tokens at position $T-\tau$, for $\tau=1,2,...$ We report this measure for the first and last layers, as well as an intermediate layer from the semantic region (layers 43 of DeepSeek-V3 and 23 of Llama3.1-8b, respectively). For both LLMs, the II rapidly grows as a function of τ for the initial layers and saturates to 1, clearly showing that the first layer is short-range correlated. For the deeper layers that were shown in Fig. 1 to correspond to the phase of shared cross-linguistic information, the II exhibits the slowest growth, meaning that far away tokens are predictive about each other, i.e., in these layers the information is mostly shared between tokens. DeepSeek's deep representations have notably lower II across tokens than Llama3.1-8b's, making it evident that better-quality semantic representations more clearly capture long-distance correlations. A clear separation in effect strength between the intermediate semantic layer and the last layer (61 for DeepSeek-V3, 32 for Llama3.1-8b) is moreover only present for DeepSeek-V3. Still, even for DeepSeek-V3 the last layer has II values clearly below those of the first layer, suggesting that, while the last layer is optimized to write the output, it still carries semantic content.

3.1.4 CORRELATIONS BETWEEN TOKENS AS A HALLMARK OF QUALITY IN SEMANTIC REPRESENTATIONS

Fig. 3a) focuses on a semantic layer from DeepSeek-V3, and it shows token-token II as a function of token distance for several languages. We present the II in both directions: The "backward" II, from the last token, T, to a previous token at distance τ , $\Delta(T \to T - \tau)$, and the reciprocal or "forward" II, $\Delta(T - \tau \to T)$. We observe that English deep representations are much more correlated than the representations in the other languages. We conjecture that this is a sign of better quality, consistent with Fig. 1b), where the presence of English in a pair results in the best predictability (lowest II scores). At the opposite extreme, note that single-language German and Hungarian representations are the least correlated in Fig. 3a), and the German-Hungarian translations have indeed the worst mutual predictability score (highest II) in Fig. 1. Furthermore, Spanish and Italian have the second and third most correlated inner representations in Fig. 3a), coherently with the profiles of the English-Spanish and English-Italian pairs in Fig. 1b) (lowest and second lowest IIs, respectively).

325

326

327

328

329

330

331

333

334 335 336

337

338

339

340

341

342 343 344

345

346

347

348

349

350

351

352

353

354

355

356

357

358

359 360

361 362

363

364

366

367

368

369

370 371

372

373

374

375

376

377

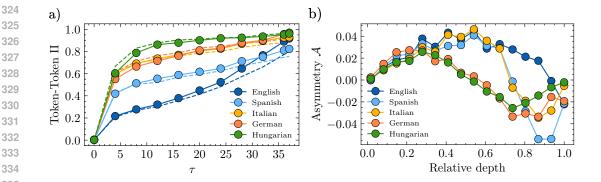


Figure 3: Panel a) In solid markers, the "backward" Information Imbalance (II), i.e., from the last token to a previous token at token-distance τ , as a function of τ , for DeepSeek-V3 representations at layer 43 and sentences in several languages. In dashed lines, the reciprocal "forward" II. Panel b) Asymmetry A between forward and backward II, averaged over τ , as a function of DeepSeek's relative depth. Error bars computed with 5 repetitions of a Jackknife procedure, subsampling 1000 sentences out of 2000, are smaller than marker size.

Finally, Fig. 3a) shows that, for English, as the distance between tokens τ increases, the relative information content between past and future tokens gets increasingly asymmetric. In particular, earlier tokens are more informative of the last one than the other way around. This could be expected, since causal models like DeepSeek-V3 are trained to predict future tokens, and text itself has a natural causal asymmetry. Nonetheless, Fig. 3a) shows that this past-to-future asymmetry is heterogeneous across languages. To further study the effect, we first define the asymmetry as $\mathcal{A} = \Delta(T \to T - \tau) - \Delta(T - \tau \to T)$. When \mathcal{A} is positive, the earlier token predicts the last one more than the latter predicts the earlier token, whereas a negative \mathcal{A} has the opposite interpretation. Fig. 3b) shows the information asymmetry A of DeepSeek-V3 as a function of its layers for the all languages we studied. For initial layers (i.e., relative depth less than roughly 0.4, see Fig. 1), we observe that all representations have a comparable amount of causal asymmetry. Instead, for semantic layers there are strong heterogeneities between languages. English remains causally asymmetric until the last few layers of the network, while the other languages present earlier negative values of A. Intriguingly, Spanish and Italian, the two languages with the lowest IIs from English, have a later dip in \mathcal{A} compared to German and Hungarian. These results call for further experiments to study possible relationships between these observed information asymmetries, model performance, and linguistic structure. We leave them as material for future research.

3.2 Images sharing high level semantics

We extend the analysis to pairs of images and measure how vision transformers (ViTs) capture semantic information for two systems: image-gpt Chen et al. (2020) and DinoV2 Oquab et al. (2024). Image-gpt performs next-token prediction after it (i) down-scales the images; (ii) unrolls them, juxtaposing the rows; (iii) quantizes the colors. It mimics the strategy of LLMs: the last layers must perform next-token prediction; previous work suggests that semantics appears in the middle Chen et al. (2020). DinoV2 encodes salient features. It trains a student network to mimic its teacher and reconstruct different instances of the same image. The final layer should be the most semantically rich: its output is the input to downstream tasks such as depth estimation, image segmentation, and instance retrieval.

We first compare representations for pairs of images from the Imagenet1k dataset Deng et al. (2009). Each couple shares the same class, which we take as a proxy of similar semantics: we process the pairs of instances with the same architecture and compute the information imbalance on a random subset of 2500 same-class pairs of images, which we sample without repetition. We replicate the procedure five times to quantify uncertainty. We report the results in the left panel of Fig. 4. As the assignment of same-pair images to two spaces used to compute II is arbitrary, we report the average of II in the two directions. The minimum of the information imbalance is in the middle for image-gpt and at the end for DinoV2, consistently with previous work Valeriani et al. (2023). The information

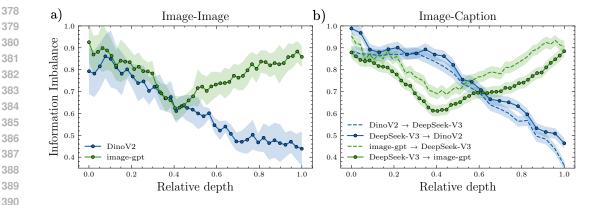


Figure 4: **Panel a**) Information imbalance for image pairs from the Imagenet1k dataset. We sample 2500 pairs of images from the same class at random; averaging over five replications. We report the mean and standard deviation. **Panel b**) Information imbalance between image and caption representations on the flickr30k dataset. Images are encoded using DinoV2 and image-gpt-large, while captions are processed with DeepSeek-V3. We report the imbalance as a function of the relative depth in the vision transformer, using the 52nd layer of DeepSeek-V3 (inside the semantic region) for captions. Results are averaged over 5000 samples, with uncertainty estimated via bootstrapping (100 replicas of 200 samples each). Dotted lines indicate information flow from caption to image; dashed lines, from image to caption.

imbalance is lower for DinoV2 (≈ 0.43): this would correspond to roughly 10% shared features in the Gaussian case, see Supp. Inf., Sec. C.

3.3 IMAGES AND CAPTIONS: MULTIMODAL DATA SHARING SEMANTIC CONTENT

Image-caption pairs encode the same semantic content in different modalities: we are interested in their representations as semantics must play a key role in alignment. We process text with DeepSeek-V3 and images with DinoV2 and image-gpt. The dataset is flickr30k Young et al. (2014): it contains pairs of images and captions; the captions are the concatenations of five different human-generated descriptions. We next repeat the analysis of Fig. 2a), estimating the minimum Information Imbalance between the image transformer and DeepSeek-V3 as a function of the number of image tokens concatenated from the end of the sequence. For DeepSeekV3, we fix the token window to 9, which corresponds to the point where the steep decrease of II observed in Fig. 2a) essentially saturates. We find that, for both DinoV2 and image-gpt, it is necessary to concatenate the last 200 tokens to capture the relevant shared information. Results obtained with shorter windows (last 50 tokens) and with mean-token are reported for comparison in Supp. Inf., Sec. J.

In Fig. 4 we report the information imbalance between layer 52 of DeepSeek-V3 (captions), DinoV2 and image-gpt (images). We select layer 52 of DeepSeek-V3 because it is a layer in which the representations of text in different languages is highly mutually predictive, see Fig. 1. We recover the exact same semantic regions provided by the analysis of the image-image pairs: towards the end for DinoV2 and in the middle for image-gpt. DinoV2 has a lower information imbalance of 0.4. Using the Gaussian process as a reference (Supp. Inf., Sec. C), this would correspond to approximately 10% of shared features. The minimum II for image-gpt is ≈ 0.6 , corresponding to 5% of shared features. The dashed lines in Fig. 4 report the information imbalance from the image to the text representation. Remarkably, for the image-gpt model the II is considerably larger than in the text-toimage direction. In particular, the minimum value of the II is ≈ 0.7 . The difference between the II in the two direction is significantly higher than the error bars estimated by bootstrap, indicating that the effect is statistically significant. This difference is likely due to the better quality of the DeepSeek-V3 representations. The asymmetry in the information imbalance is reversed for the DinoV2 model. In this case the image representation predicts the text representation marginally better than the reverse. This can be due to the fact that DinoV2 is explicitly trained to extract semantically meaningful visual features, and indeed the largest asymmetry is observed in the last layer.

To directly test whether joint text–image training facilitates cross-modal alignment, we performed a set of experiments with the ViT-L/14 visual encoder of the multimodal CLIP system (Radford et al., 2021) on the same flickr30k dataset (Supp. Inf., Sec. J). We computed the information imbalance (II) between CLIP image representations and DeepSeek-V3 text representations and compared it with the DinoV2–DeepSeek-V3 baseline. We found that the DeepSeekV3-CLIP Image II reaches a minimum of approximately 0.3 ($\approx 15\%$ of shared features), which is lower than the minimum II observed for DinoV3-DeepSeeV3, indicating that CLIP's visual encoder aligns more closely to the text semantics than a purely visual model of comparable capacity.

4 LIMITATIONS

It would be interesting to extend our analyses to a further variety of models of intermediate sizes and diverse architectures, including diffusion models, for example quantifying the depth and the length of the semantic regions found as a function of the number of parameters in the model, and to consider models that specifically target multilinguality Martins et al. (2025); Üstün et al. (2024). In our experiments with translations, we used human translations from opus_books Tiedemann (2012), in which the different language pairs may correspond to translations from different novels, introducing an extra source of variation not taken into account in our analysis. We explicitly avoided using LLMs to generate the translations to avoid the introduction of biases in the analyses. Indeed, we found that, for the case of English-Italian, DeepSeek-generated translations have more similar representations than those obtained from opus_books. More generally, in our multilingual analyses of DeepSeek-V3, there are more comparisons to be made, such as fixing English as a pivot language and then systematically considering different translations of the same English text into languages with different online presence, and of different degree of relatedness.

5 CONCLUSION

Through a novel application of the information imbalance measure Glielmo et al. (2022), we showed that DeepSeek-V3, the largest publicly available LLM, developed an internal processing phase in which different inputs that share the same semantics, such as translations and image-caption pairs, are reflected in representations that are extremely similar. We further ascertained that, on these "semantic" layers, long token spans meaningfully contribute to the representation, and long-distance correlations in encoded information cue high-quality representations. Several of the patterns we observed in DeepSeek-V3 also emerge in the medium-sized Llama3.1-8b model, but in a less distinct manner, suggesting that deep semantic processing is a hallmark of better models. We also analyzed visual processing models, finding again the presence of layers that capture deeper semantic similarity, and whose position depends on the objective the models are trained upon. DeepSeek-V3 textual representations and vision transformer visual representations of the same concepts are most strongly aligned in the respective semantic layers.

Our work supports the hypothesis that, as deep models improve, they converge towards shared representations of the world Huh et al. (2024). We took a first step away from simply verifying that such representations exist, towards characterizing their nature, by precisely quantifying the degree of shared information, localizing where they occur in different networks and determining how they are synthesized from composite inputs. Future research, besides extending the empirical coverage of our investigation, should further increase the resolution of the measurements, allowing us to go from observing the holistic behavior of a model with respect to a data set to a full characterization of how semantic representations are constructed given each specific input instance.

6 ETHICS STATEMENT

This work adheres to the ICLR Code of Ethics. Our experiments employ publicly available datasets and open-source models (under their respective licenses). We are not aware of any ethical concerns. To the contrary, by providing new tools to understand the inner workings of LLMs and other AI systems, we hope to contribute to a safer and more transparent AI.

7 REPRODUCIBILITY STATEMENTS

We describe all model configurations, datasets, and computational settings in the main text and supplementary material to enable independent verification. Upon acceptance, we will release the full source code and instructions in a public repository to facilitate reproducibility.

REFERENCES

486

487 488

489

490

491 492

493 494

495

496

497 498

499

500

501

502

504

505

506

507

508

509 510

511

512

513

514

515516

517

518 519

520

521

522

523

524

525

526

527

528

529

530

531

532

534

536

538

- Languages Used on the Internet // Wikipedia. IV 2025.
- Acevedo Santiago, Rodriguez Alex, Laio Alessandro. Unsupervised detection of semantic correlations in big data // Communications Physics. 2025. 8, 1. 202.
- Anderson Alexander G., Berg Cory P. The High-Dimensional Geometry of Binary Neural Networks // International Conference on Learning Representations. 2018.
- Antonello Richard, Cheng Emily. Evidence from fMRI Supports a Two-Phase Abstraction Process in Language Models // UniReps: 2nd Edition of the Workshop on Unifying Representations in Neural Models. 2024.
- Bansal Yamini, Nakkiran Preetum, Barak Boaz. Revisiting Model Stitching to Compare Neural Representations // Advances in Neural Information Processing Systems. 34. 2021. 225–236.
- Bondarenko Yelysei, Nagel Markus, Blankevoort Tijmen. Quantizable Transformers: Removing Outliers by Helping Attention Heads Do Nothing // Thirty-seventh Conference on Neural Information Processing Systems. 2023.
- Brinkmann Jannik, Wendler Chris, Bartelt Christian, Mueller Aaron. Large Language Models Share Representations of Latent Grammatical Concepts Across Typologically Diverse Languages. 2025.
- Chen Mark, Radford Alec, Child Rewon, Wu Jeffrey, Jun Heewoo, Luan David, Sutskever Ilya. Generative Pretraining From Pixels // Proceedings of the 37th International Conference on Machine Learning. 119. 13–18 Jul 2020. 1691–1703. (Proceedings of Machine Learning Research).
- Cheng Emily, Doimo Diego, Kervadec Corentin, Macocco Iuri, Yu Lei, Laio Alessandro, Baroni Marco. Emergence of a High-Dimensional Abstraction Phase in Language Transformers // The Thirteenth International Conference on Learning Representations. 2025.
- DeepSeek-AI, Liu Aixin, Feng Bei, Xue Bing, Wang Bingxuan, Wu Bochao, Lu Chengda, Zhao Chenggang, Deng Chengai, Zhang Chenyu, Ruan Chong, Dai Damai, Guo Daya, Yang Dejian, Chen Deli, Ji Dongjie, Li Erhang, Lin Fangyun, Dai Fucong, Luo Fuli, Hao Guangbo, Chen Guanting, Li Guowei, Zhang H., Bao Han, Xu Hanwei, Wang Haocheng, Zhang Haowei, Ding Honghui, Xin Huajian, Gao Huazuo, Li Hui, Qu Hui, Cai J. L., Liang Jian, Guo Jianzhong, Ni Jiaqi, Li Jiashi, Wang Jiawei, Chen Jin, Chen Jingchang, Yuan Jingyang, Qiu Junjie, Li Junlong, Song Junxiao, Dong Kai, Hu Kai, Gao Kaige, Guan Kang, Huang Kexin, Yu Kuai, Wang Lean, Zhang Lecong, Xu Lei, Xia Leyi, Zhao Liang, Wang Litong, Zhang Liyue, Li Meng, Wang Miaojun, Zhang Mingchuan, Zhang Minghua, Tang Minghui, Li Mingming, Tian Ning, Huang Panpan, Wang Peiyi, Zhang Peng, Wang Qiancheng, Zhu Qihao, Chen Qinyu, Du Qiushi, Chen R. J., Jin R. L., Ge Ruiqi, Zhang Ruisong, Pan Ruizhe, Wang Runji, Xu Runxin, Zhang Ruoyu, Chen Ruyi, Li S. S., Lu Shanghao, Zhou Shangyan, Chen Shanhuang, Wu Shaoqing, Ye Shengfeng, Ye Shengfeng, Ma Shirong, Wang Shiyu, Zhou Shuang, Yu Shuiping, Zhou Shunfeng, Pan Shuting, Wang T., Yun Tao, Pei Tian, Sun Tianyu, Xiao W. L., Zeng Wangding, Zhao Wanjia, An Wei, Liu Wen, Liang Wenfeng, Gao Wenjun, Yu Wenqin, Zhang Wentao, Li X. Q., Jin Xiangyue, Wang Xianzu, Bi Xiao, Liu Xiaodong, Wang Xiaohan, Shen Xiaojin, Chen Xiaokang, Zhang Xiaokang, Chen Xiaosha, Nie Xiaotao, Sun Xiaowen, Wang Xiaoxiang, Cheng Xin, Liu Xin, Xie Xin, Liu Xingchao, Yu Xingkai, Song Xinnan, Shan Xinxia, Zhou Xinyi, Yang Xinyu, Li Xinyuan, Su Xuecheng, Lin Xuheng, Li Y. K., Wang Y. Q., Wei Y. X., Zhu Y. X., Zhang Yang, Xu Yanhong, Xu Yanhong, Huang Yanping, Li Yao, Zhao Yao, Sun Yaofeng, Li Yaohui, Wang Yaohui, Yu Yi, Zheng Yi, Zhang Yichao, Shi Yifan, Xiong Yiliang, He Ying, Tang Ying, Piao Yishi, Wang Yisong, Tan Yixuan, Ma Yiyang, Liu Yiyuan, Guo Yongqiang, Wu Yu, Ou Yuan, Zhu Yuchen, Wang Yuduan, Gong Yue, Zou Yuheng, He Yujia, Zha Yukun, Xiong Yunfan, Ma Yunxian, Yan Yuting, Luo Yuxiang, You Yuxiang, Liu Yuxuan, Zhou Yuyang,

Wu Z. F., Ren Z. Z., Ren Zehui, Sha Zhangli, Fu Zhe, Xu Zhean, Huang Zhen, Zhang Zhen, Xie
 Zhenda, Zhang Zhengyan, Hao Zhewen, Gou Zhibin, Ma Zhicheng, Yan Zhigang, Shao Zhihong,
 Xu Zhipeng, Wu Zhiyu, Zhang Zhongyu, Li Zhuoshu, Gu Zihui, Zhu Zijia, Liu Zijun, Li Zilin, Xie
 Ziwei, Song Ziyang, Gao Ziyi, Pan Zizheng. DeepSeek-V3 Technical Report. 2025.

- Deng Jia, Dong Wei, Socher Richard, Li Lia-Ji, Fei-Fei Li. ImageNet: A Large-Scale Hierarchical Image Database // Proceedings of CVPR. Miami Beach, FL, 2009. 248–255.
- Doimo Diego, Glielmo Aldo, Ansuini Alessio, Laio Alessandro. Hierarchical nucleation in deep neural networks // Proceedings of NeurIPS. 2020. Published online: https://papers.nips.cc/paper/2020.
- Glielmo Aldo, Zeni Claudio, Cheng Bingqing, Csányi Gábor, Laio Alessandro. Ranking the information content of distance measures // PNAS Nexus. 04 2022. 1, 2. pgac039.
- Guo Yunhui. A survey on methods and theories of quantized neural networks // arXiv preprint arXiv:1808.04752. 2018.
- Hubara Itay, Courbariaux Matthieu, Soudry Daniel, El-Yaniv Ran, Bengio Yoshua. Binarized Neural Networks // Advances in Neural Information Processing Systems. 29. 2016.
- Hubara Itay, Courbariaux Matthieu, Soudry Daniel, El-Yaniv Ran, Bengio Yoshua. Quantized Neural Networks: Training Neural Networks with Low Precision Weights and Activations // Journal of Machine Learning Research. 2018. 18, 187. 1–30.
- Huh Minyoung, Cheung Brian, Wang Tongzhou, Isola Phillip. Position: The Platonic Representation Hypothesis // Proceedings of the 41st International Conference on Machine Learning. 235. 21–27 Jul 2024. 20617–20642. (Proceedings of Machine Learning Research).
- Kornblith Simon, Norouzi Mohammad, Lee Honglak, Hinton Geoffrey. Similarity of Neural Network Representations Revisited // Proceedings of ICML. Long Beach, CA, 2019. 3519–3529.
- Kovaleva Olga, Kulshreshtha Saurabh, Rogers Anna, Rumshisky Anna. BERT Busters: Outlier Dimensions that Disrupt Transformers // Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021. Online: Association for Computational Linguistics, VIII 2021. 3392–3405.
- Lenc Karel, Vedaldi Andrea. Understanding Image Representations by Measuring Their Equivariance and Equivalence // International Journal of Computer Vision. 05 2019. 127.
- Lindsey Jack, Gurnee Wes, Ameisen Emmanuel, Chen Brian, Pearce Adam, Turner Nicholas L., Citro Craig, Abrahams David, Carter Shan, Hosmer Basil, Marcus Jonathan, Sklar Michael, Templeton Adly, Bricken Trenton, McDougall Callum, Cunningham Hoagy, Henighan Thomas, Jermyn Adam, Jones Andy, Persic Andrew, Qi Zhenyi, Thompson T. Ben, Zimmerman Sam, Rivoire Kelley, Conerly Thomas, Olah Chris, Batson Joshua. On the Biology of a Large Language Model // Transformer Circuits Thread. 2025.
- Maniparambil Mayug, Akshulakov Raiymbek, Djilali Yasser Abdelaziz Dahou, Narayan Sanath, Seddik Mohamed El Amine, Mangalam Karttikeya, O'Connor Noel. Do vision and language encoders represent the world similarly? // Proceedings of CVPR. Seattle, WA, 2024. 14334–14343.
- Martins Pedro Henrique, Fernandes Patrick, Alves João, Guerreiro Nuno M., Rei Ricardo, Alves Duarte M., Pombal José, Farajian Amin, Faysse Manuel, Klimaszewski Mateusz, Colombo Pierre, Haddow Barry, De Souza José G.C., Birch Alexandra, Martins André F.T. EuroLLM: Multilingual Language Models for Europe // Procedia Computer Science. 2025. 255. 53–62. Publisher Copyright: © 2025 The Authors. Published by Elsevier B.V.; 2nd EuroHPC user day, EuroHPC 2024; Conference date: 22-10-2024 Through 23-10-2024.
- Meta. Introducing Meta Llama 3: The most capable openly available LLM to date. 2024.
- Moschella Luca, Maiorca Valentino, Fumero Marco, Norelli Antonio, Locatello Francesco, Rodolà Emanuele. Relative representations enable zero-shot latent space communication // The Eleventh International Conference on Learning Representations. 2023.

Oquab Maxime, Darcet Timothée, Moutakanni Théo, Vo Huy, Szafraniec Marc, Khalidov Vasil, Fernandez Pierre, Haziza Daniel, Massa Francisco, El-Nouby Alaaeldin, Assran Mahmoud, Ballas Nicolas, Galuba Wojciech, Howes Russell, Huang Po-Yao, Li Shang-Wen, Misra Ishan, Rabbat Michael, Sharma Vasu, Synnaeve Gabriel, Xu Hu, Jegou Hervé, Mairal Julien, Labatut Patrick, Joulin Armand, Bojanowski Piotr. DINOv2: Learning Robust Visual Features without Supervision. II 2024.

- Peng Qiwei, Søgaard Anders. Concept Space Alignment in Multilingual LLMs. 2024.
- Radford Alec, Kim Jong Wook, Hallacy Chris, Ramesh Aditya, Goh Gabriel, Agarwal Sandhini, Sastry Girish, Askell Amanda, Mishkin Pamela, Clark Jack, others. Learning transferable visual models from natural language supervision // International conference on machine learning. 2021. 8748–8763.
- Singh Jasdeep, McCann Bryan, Socher Richard, Xiong Caiming. BERT is Not an Interlingua and the Bias of Tokenization // Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019). Hong Kong, China: Association for Computational Linguistics, XI 2019. 47–55.
- Sorscher Ben, Ganguli Surya, Sompolinsky Haim. Neural representational geometry underlies few-shot concept learning // Proceedings of the National Academy of Sciences. 2022. 119, 43. e2200800119.
- Sucholutsky Ilia, Muttenthaler Lukas, Weller Adrian, Peng Andi, Bobu Andreea, Kim Been, Love Bradley C., Cueva Christopher J., Grant Erin, Groen Iris, Achterberg Jascha, Tenenbaum Joshua B., Collins Katherine M., Hermann Katherine L., Oktar Kerem, Greff Klaus, Hebart Martin N., Cloos Nathan, Kriegeskorte Nikolaus, Jacoby Nori, Zhang Qiuyi, Marjieh Raja, Geirhos Robert, Chen Sherol, Kornblith Simon, Rane Sunayana, Konkle Talia, O'Connell Thomas P., Unterthiner Thomas, Lampinen Andrew K., Müller Klaus-Robert, Toneva Mariya, Griffiths Thomas L. Getting aligned on representational alignment. 2024.
- Sun Mingjie, Chen Xinlei, Kolter J Zico, Liu Zhuang. Massive Activations in Large Language Models // First Conference on Language Modeling. 2024.
- Tatto Vittorio Del, Fortunato Gianfranco, Bueti Domenica, Laio Alessandro. Robust inference of causality in high-dimensional dynamical processes from the Information Imbalance of distance ranks // Proceedings of the National Academy of Sciences. 2024. 121, 19. e2317256121.
- *Tiedemann Jörg.* Parallel Data, Tools and Interfaces in OPUS // Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12). Istanbul, Turkey: European Language Resources Association (ELRA), V 2012. 2214–2218.
- Üstün Ahmet, Aryabumi Viraat, Yong Zheng, Ko Wei-Yin, D'souza Daniel, Onilude Gbemileke, Bhandari Neel, Singh Shivalika, Ooi Hui-Lee, Kayid Amr, Vargus Freddie, Blunsom Phil, Longpre Shayne, Muennighoff Niklas, Fadaee Marzieh, Kreutzer Julia, Hooker Sara. Aya Model: An Instruction Finetuned Open-Access Multilingual Language Model // Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). Bangkok, Thailand: Association for Computational Linguistics, VIII 2024. 15894–15939.
- Valeriani Lucrezia, Doimo Diego, Cuturello Francesca, Laio Alessandro, Ansuini Alessio, Cazzaniga Alberto. The geometry of hidden representations of large transformer models // Advances in Neural Information Processing Systems. 2023. 36. 51234–51252.
- Wang Hongyu, Ma Shuming, Dong Li, Huang Shaohan, Wang Huaijie, Ma Lingxiao, Yang Fan, Wang Ruiping, Wu Yi, Wei Furu. Bitnet: Scaling 1-bit transformers for large language models // arXiv preprint arXiv:2310.11453. 2023.
- Wendler Chris, Veselovsky Veniamin, Monea Giovanni, West Robert. Do Llamas Work in English? On the Latent Language of Multilingual Transformers // Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). Bangkok, Thailand: Association for Computational Linguistics, VIII 2024. 15366–15394.

Wild Romina, Wodaczek Felix, Del Tatto Vittorio, Cheng Bingqing, Laio Alessandro. Automatic Feature Selection and Weighting in Molecular Systems Using Differentiable Information Imbalance // Nature Communications. I 2025. 16, 1. 270.

- *Young Peter, Lai Alice, Hodosh Micah, Hockenmaier Julia.* From Image Descriptions to Visual Denotations: New Similarity Metrics for Semantic Inference over Event Descriptions // Transactions of the Association for Computational Linguistics. XII 2014. 2. 67–78.
- Zhao Yiran, Zhang Wenxuan, Chen Guizhen, Kawaguchi Kenji, Bing Lidong. How do Large Language Models Handle Multilingualism? // Advances in Neural Information Processing Systems. 37. 2024. 15296–15319.
- Zheng Lianmin, Yin Liangsheng, Xie Zhiqiang, Sun Chuyue, Huang Jeff, Yu Cody Hao, Cao Shiyi, Kozyrakis Christos, Stoica Ion, Gonzalez Joseph E., Barrett Clark, Sheng Ying. SGLang: Efficient Execution of Structured Language Model Programs // The Thirty-eighth Annual Conference on Neural Information Processing Systems. 2024.

Supplementary Information

A COMPUTE RESOURCES

We run DeepSeek-V3 on a cluster of 16 H100 GPUs (80GB each), using the SGLang framework Zheng et al. (2024). Input sequences consist of 40–80 tokens and are processed in batches of 30. For LLama3.1-8B, a single GPU from a DGX H100 system is sufficient. For our experiments on text and image-caption pairs, 1TB of RAM is enough, moving to the GPU a couple of layers at a time. Image-Image experiments are done using up to five NVIDIA A30 GPUs and 500 GB of RAM.

B ASSETS

Opus books: https://huggingface.co/datasets/Helsinki-NLP/opus_books; licence: CC BY 4.0.

 $\label{eq:DeepSeek-V3:Model licence} DeepSeek-V3: \ https://huggingface.co/deepseek-ai/DeepSeek-V3: \ Model \ licence.$

Llama3.1-8b: https://huggingface.co/meta-llama/Llama3.1-8b-3.1-8B; license: llama3.1

Flickr dataset: https://huggingface.co/datasets/clip-benchmark/wds_flickr30k; subject to flickr terms of use.

Imagenet: https://image-net.org/; terms of access: https://image-net.org/download.php

DinoV2: https://github.com/facebookresearch/dinov2/tree/main; licence: Apache 2.0

image-gpt-large: https://github.com/openai/image-gpt; licence: Modified MIT

C Information Imbalance: implementation details and Gaussian comparisons

Implementation details The representations of deep transformers are (very) high dimensional: the number of tokens ranges from a few tens for short sentences, to thousands for long sentences or images, and typical (per token) embedding dimensions are of the order $\mathcal{O}(10^3)$. These two contributions rapidly generate representations where the total dimensionality can be of order $\mathcal{O}(10^6)$. Estimating the information imbalance in such a large space can be very costly, especially with sample sizes of order 10^3 , like those used in this work. To speed up computations, we follow the approach introduced in Ref. Wild et al. (2025). Instead of directly computing all-to-all distances between the N samples, we randomly split the data, selecting two subsets of points, that we call origins and destinations. We call N_o and N_d the number of samples in both subsets, respectively. Computing the information imbalance without splitting is costly for two reasons. First, it is necessary to compute the pairwise distances between all N data points, instead of using only the N_0 origins. Second, the ranking operations sort N-1 elements instead of N_d-1 elements, which strongly affects compute time, as sorting operations grow more than linearly with the number of elements. In this work, we found a good number of origin points to be $\approx 200 - 500$. Indeed, with these numbers the statistical uncertainty on the II is small enough to assess subtle differences between the languages, as shown in the figures in the main text (notice that the error bars are smaller than the marker size in all the figures except in Fig. 4, where the error bars are explicitly shown).

Interpreting II To provide an interpretation of the value of the information imbalance observed in our analysis of the representations of images and text, we computed the II in an idealized scenario, in which the dimensions of the feature spaces resemble the one in neural networks: we generate set Y with 2000 samples from a Gaussian distribution with zero-mean, identity variance matrix, and $p=10^5$ components. We fix X to be the first k components of Y, with $k=1,2,\ldots,10^5$. We compute the information imbalance $\Delta(X\to Y)$ to quantify how adding features changes the II. We repeat the analysis by binarizing the features, namely following the same procedure used for the activations in the neural networks. In the real case, we use the L^2 distance; in the binary case we employ the Hamming distance. We report the results in Fig.5: we average the results over 10 replicas.

We see that the information imbalance between X and Y decreases as the number of shared features increases in a highly nonlinear manner, and that the results with and without binarization are within statistical error. In Sec. H we further show that binarization also preserves the Information Imbalance calculated with activations generated by DeepSeek-V3.

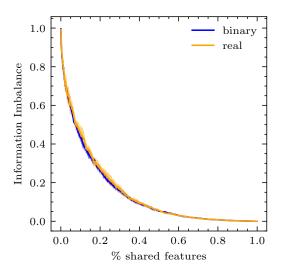


Figure 5: Information imbalance for growing shared components in simulated data. We compute the information imbalance $\Delta(X \to Y)$, where Y is the target space and X is the origin space: it consists of the first k features of Y, with $k=1,\ldots,10^5$. We report both the real (Euclidean distance) and the binary cases (Hamming distance). The results are averaged across ten repetitions; the shades are the 1-sigma bootstrap errors.

D TRANSLATIONS DATASET

We filter pairs of translated sentences from Helsinki-NLP/opus_books Tiedemann (2012) that have a number of tokens between 40 and 80, avoiding trivial sentences and incorrect translations, while keeping roughly 5000 thousands sentences of comparable length. We avoid truncating the translated sentences, preserving their full meaning. We excluded the last two tokens of each sentence from the analysis, as they consistently correspond to punctuation marks (e.g., a period or a period followed by a quotation mark), which introduce trivial similarities in the representations.

E AVERAGE TOKEN COMPARISON

Fig. 6 shows the Information Imbalance between DeepSeek-V3 representations of English and Spanish translations, using 20 tokens of each sentence, with and without taking the average over the token axis. Note that for the case of concatenated tokens, the curve is the same as in Fig. 1 of the main article, for direct comparison. When taking the average, we use the L_2 distance and we clip and normalize the inputs, as Ref. Huh et al. (2024). When using all 20 tokens we binarize activations and we use the Hamming distance, as in the main text. While the results are qualitatively consistent, they differ significantly in quantitative terms. Across all layers—most notably in the first and last—averaging leads to a marked increase in mutual predictability (lower Information Imbalance), due to the elimination of positional information, which is most relevant in the initial and last layers, i.e., before and after the semantic region.

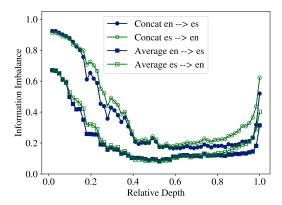


Figure 6: Information Imbalance between English (en) and Spanish (es) representations generated by DeepSeek-V3's as a function of depth. 'Concat' stands for the results obtain with the concatenation of the 20 last tokens of each sentence. 'Average' stands for the average of the same 20 tokens. The standard deviation is computed with a Jackknife procedure, subsampling 2000 samples out of 5000 five times, and it is smaller than marker size.

F MISALIGNMENT OF TRANSLATIONS ERASES SEMANTIC SIMILARITY

As a consistency check, Fig 7 shows the Information Imbalance for DeepSeek-V3's and Llama3.1-8b representations using misaligned translations, namely performing a batch-shuffle in one of the datasets. Since the semantic correspondence between sentences is destroyed, the representations are not informative about each other, and thus the II is close to one, for all layers. The same occurs with misaligned image pairs, II being around one when pairs do not share the same class, for all layers.

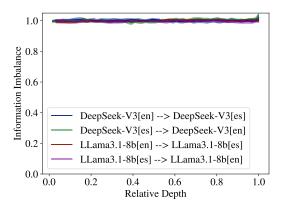


Figure 7: Information Imbalance between English (en) and Spanish (es) representations generated by DeepSeek-V3 and Llama3.1-8b, for the "non-informative case", namely a misaligned dataset in which we batch-shuffle the Spanish translations. The (hardly visible) shaded area corresponds to one standard deviation, computed with a Jackknife procedure subsampling 2000 samples out of 5000.

G COMPARISON BETWEEN NEIGHBORHOOD OVERLAP AND INFORMATION IMBALANCE

To quantify representation similarity, Ref. Huh et al. (2024) used the neighborhood overlap, namely the average fraction of shared k nearest neighbors. In particular, they found this metric to be more suitable than linear metrics, like Central Kernel Alignment (CKA), providing a stronger signal. In this section we compare our results from Fig. 1.a) of the main article with the neighborhood overlap computed on the same representations. To measure the neighborhood overlap we follow the pipeline of Ref. Huh et al. (2024), namely we clip the activations using quantiles of order 0.05% and 95%, we

normalize each vector, and we measure the distances between them with the L_2 metric. Fig 8 shows that the neighborhood overlap between representations increases and reaches a plateau concurrently with the local minimum plateau of the Information Imbalance, thus both methods give qualitative the same results. We note that the neighborhood overlap, being defined between 0 and 1, here only reaches a value around 0.2. Similar small values were also observed and even highlighted in Ref. Huh et al. (2024), although the authors do not provide arguments to understand the reasons behind this property. We note that if the k-th neighbor of representation A is instead the k+1 neighbor of representation B, automatically that point counts as outside the k-neighborhood, even if it is extremely close, rendering the alignment value very low. This effect doesn't take place in the Information Imbalance, which, instead of evaluating if k neighbors are shared, measures what is the rank in space B of k neighbors in space A, and vice-versa.

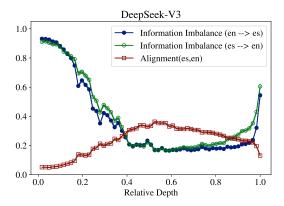


Figure 8: Information Imbalance between English (en) and Spanish (es) representations generated by DeepSeek-V3 from Fig. 1 of the main article, compared with the neighborhood alignment metric from Ref. Huh et al. (2024). The (hardly visible) shaded area corresponds to one standard deviation, computed with a Jackknife procedure subsampling 2000 samples out of 5000.

H BINARIZING HIGH DIMENSIONAL REPRESENTATIONS HAS A MARGINAL EFFECT ON THE INFORMATION IMBALANCE

Fig. 9 includes again, as reference, the Information Imbalance between DeepSeek-V3 representations of English and Spanish translations, using the last 20 tokens of each sentence, binarizing the activations and using the Hamming distance, (results of Fig. 1 from main text). To investigate the effects of binarizing the representations, we show the Information Imbalance using the same data where, instead of taking the sign of the activations to generate binary variables, we clip them using quantiles of order 0.05% and 95%, we normalize vectors to unit norm, and we use the L_2 distance, similar to Ref. Huh et al. (2024). Remarkably, binarizing has marginal effects on the Information Imbalance computed with full precision (for DeepSeek-V3's representations, BF16), which is possible in this setup given that we are working with 20 tokens.

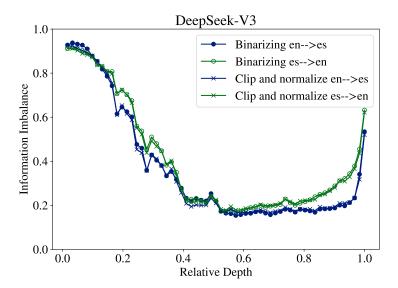


Figure 9: In circles, Information Imbalance between English (en) and Spanish (es) representations generated by DeepSeek-V3, binarizing the activations and using the Hamming distance. In crosses, the II clipping and normalizing activations, and using the L_2 metric. The standard deviation, computed with a Jackknife procedure subsampling 2000 samples out of 5000, is smaller than marker size.

I BINARY INTRINSIC DIMENSION OF DEEPSEEK-V3 REPRESENTATIONS

The Binary Intrinsic Dimension (BID) is defined in Ref. Acevedo et al. (2025) through the probability distribution of Hamming distances between samples of binary high dimensional vectors. If we consider N-dimensional vectors σ with components σ_i , i=1,...,N uniformly distributed in $\{0,1\}$, then the probability of observing Hamming distance r between any two samples is exactly $P_0(r) = \frac{1}{2^N} \binom{N}{r}$. The BID is defined as the coefficient d_0 of the following Ansatz:

$$P(r) = \frac{1}{2^{(d_0 + d_1 r)}} \binom{N}{d_0 + d_1 r},\tag{2}$$

where d_1 is a second variational parameter. We perform a second-order optimization of the Kullback-Leibler divergence between the model equation 2 and the empirical distribution of distances, computing numerical derivatives with JAX's autodiff. For more details on the BID computation, see Acevedo et al. (2025). We used hyperparameters $\alpha_{min}=0.15$, and $\alpha_{max}=0.4$. Fig. 10 shows the Binary Intrinsic Dimension (BID) of DeepSeek-V3's representations processing English sentences from opus_books with length between 40 and 50 tokens. We take the last 10 and the last 20 tokens of the binarized representation, and we find two dimensionality peaks that roughly coincide with the beginning (around 0.4) and the end (around 0.9) of the semantic region found in Fig. 1. These results align with recent evidence of abstract spaces presenting dimensionality peaks in deep representations Cheng et al. (2025); Antonello, Cheng (2024). The missing points at the beginning of the blue curve correspond to failed optimizations due to the strong multimodality of the empirical distribution of distances.

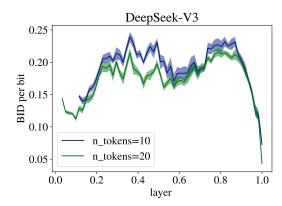


Figure 10: Binary Intrinsic Dimension (BID) normalized by the number of bits (neurons), in the representation generated by DeepSeek-V3 processing English text. In the legend, n_tokens stands for the number of tokens concatenated in the representation, starting from the last. Shaded area corresponds to one standard deviation, computed with a Jackknife procedure, subsampling 2000 samples out of 5000, ten times.

J IMAGE-CAPTIONS PAIRS

In Fig. 11, we report additional results for the image-caption experiments: we consider different windows of tokens, using the last 200 tokens for images - as we do in the main text - vs a smaller windows of 50 tokens and the average. In general, results are qualitatively similar across setups. We observe that the choice of the number of tokens comes with a trade-off: more tokens convey more information; at the same time, having more tokens implies working with bigger objects.

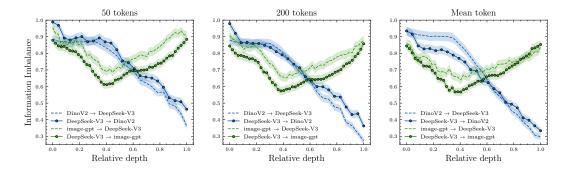


Figure 11: Information imbalance for image-caption pairs. We consider three scenarios: (i) 50 tokens; (ii) 200 tokens; and (iii) the mean token. The results are qualitatively similar.

In Fig. 12, we show how the alignment between models changes when the number of tokens is increased: on the left, we report the case for DinoV2, comparing it with DeepSeek-V3, in the image-caption case, and with itself for the pairs of images. On the right, we report the same analysis for image-gpt.

Finally, we performed three analyses on the flickr30k image-caption pairs to probe how CLIP's joint training influences representational alignment.

- 1. DinoV2 → CLIP-Image: comparing a purely visual model with CLIP's visual branch.
- 2. DeepSeek-V3 → CLIP-Text: comparing a purely textual model with CLIP's text branch.
- 3. DeepSeek-V3 → CLIP-Image: comparing a purely textual model with CLIP's visual branch.

Results are in Fig. 13. In all cases we compute the information imbalance (II) in both directions (from the source model to the target and vice-versa), always indicating the direction in the figure legends. The DeepSeek-V3 \rightarrow CLIP-Image comparison directly tests whether a visual encoder trained to align with text predicts textual representations better than a unimodal vision model. Its minimum II is approximately $0.3 \ (\approx 15\% \ \text{shared features})$, which is lower than the $\approx 0.40 \ \text{minimum}$ for DeepSeek-V3 \rightarrow DinoV2 (main text Fig. 4). The DinoV2 \rightarrow CLIP-Image II remains extremely low (≈ 0.01) past a relative depth of ≈ 0.8 , consistent with a near-identical visual representation. Conversely, the DeepSeek-V3 \rightarrow CLIP-Text II decreases more gradually and stabilizes around 0.4, indicating weaker alignment of CLIP's text branch to purely textual models. Shaded regions in Fig. 13 show 95% bootstrap confidence intervals.

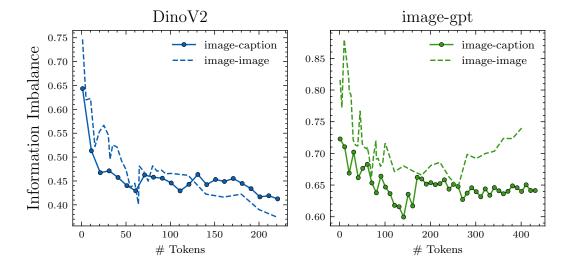


Figure 12: Information imbalance for varying windows of tokens. We report the minimum of the information imbalance, a proxy of alignment, for growing windows of tokens. On the left, we report the results for DinoV2 in both the image-caption (dotted blue line) and the image-image case (dashed blue line). On the right, we report the results for the image-gpt in both the image-caption (dotted green line) and the image-image cases (dashed green line).

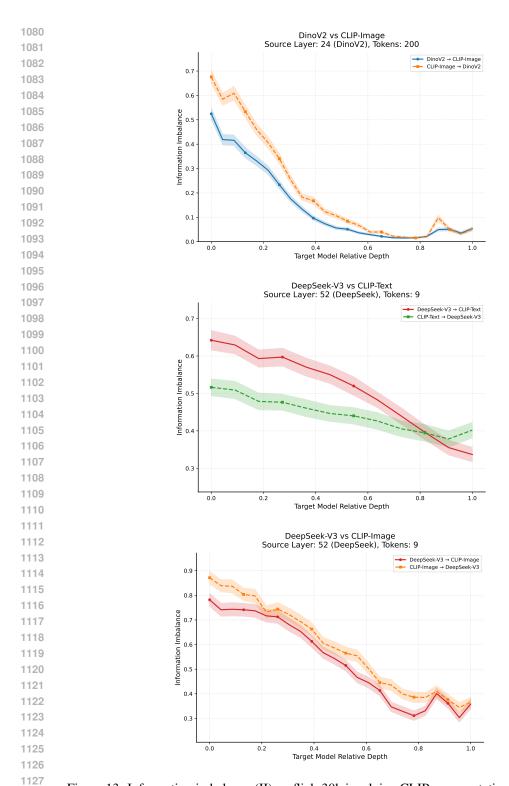


Figure 13: Information imbalance (II) on flickr30k involving CLIP representations. Top: DinoV2 image vs. CLIP image. Middle: DeepSeek-V3 text (layer 52) vs. CLIP text. Bottom: DeepSeek-V3 text vs. CLIP image (new cross-modal analysis). Lower II indicates stronger semantic predictivity. Shaded areas show 95% bootstrap confidence intervals.