

ON THE ENTROPY OF LANGUAGE MODELS IN GETTING SEMANTIC FROM TOKENS

Anonymous authors

Paper under double-blind review

ABSTRACT

Large language models (LLMs) are widely recognized for their exceptional capacity to capture semantic meaning. Yet, there remains no established metric to quantify this capability. In this work, we introduce a quantitative metric, **Information Emergence (IE)**, designed to measure LLMs’ ability to extract semantics from input tokens. We formalize “semantics” as the meaningful information abstracted from a sequence of tokens and, leveraging information theory, quantify this through comparing the reduction in entropy observed for a sequence of tokens (macro-level) and individual tokens (micro-level). To achieve this, we design a light-weight estimator to compute the mutual information at both micro and macro levels for each transformer layer, which is agnostic to different tasks and language model architectures. We apply IE in both synthetic in-context learning (ICL) scenarios and natural sentence contexts. Experiments show a high-level informativeness of our metric reflected in semantic faithfulness, sensitivity and connection with emergence. In addition, we highlight some interesting findings: 1) IE explains why ICL offers clearer semantics and benefits compared to natural text through changes in entropy. 2) We could associate certain hallucination phenomenon with increased variance in IE. 3) IE can effectively differentiate between human-written and LLM-generated text, proving especially useful for extremely large and closed-source language models. Our codes are available at: <https://anonymous.4open.science/r/Emergence/>.

1 INTRODUCTION

One of the most elusive and captivating attributes of large language models (LLMs) is their ability to learn semantics from inputs across diverse domains (Chen, 2023; Chang et al., 2024; Minaee et al., 2024), a feature that owes much to a cross-pollination of unsupervised training and next-token prediction (NTP) mechanisms. It has stimulated numerous significant research directions, such as in-context learning (ICL) (Min et al., 2022a;b; Wies et al., 2024; Ye et al., 2023; Kossen et al., 2023; Swaminathan et al., 2024), emergence capabilities (Wei et al., 2022; Schaeffer et al., 2023; Srivastava et al., 2023; Lu et al., 2023; Yu & Dong, 2022; Liu et al., 2024), and hallucination investigations (Rawte et al., 2023; Ji et al., 2023; Zhang et al., 2023; Deemter, 2024).

However, the capability of LLMs to capture semantics from texts is challenging to quantify and thereby, to evaluate. Numerous existing tasks indirectly reflect similar capabilities through evaluating LLMs’ performances (e.g., accuracy) on a specific task, such as “instruction following” (Zeng et al., 2023), “searching” (Sun et al., 2023), and “reasoning” (Yang et al., 2024). Nevertheless, these evaluation methods rely on manually curating datasets and tasks tailoring different aspects, resulting in time-consuming and domain-specific findings. In addition, these evaluations typically focus on coarse-grained text, not providing interpretations for the behavior of finer-grained tokens. Lastly, existing evaluation metrics which vary across different tasks can lead to varied performances, and even contradicting conclusions (Schaeffer et al., 2023).

In response to the above limitations, we propose a task-agnostic and closed-form metric, which we refer to as **Information Emergence (IE)**¹, designed to reflect and deterministically quantify the

¹The “Information Emergence” here and “Emergence” in LLM-related research are two different notions, we discuss their difference in Section 2.1 and Section 4.3.

ability of LLMs to extract meaningful semantics from input tokens. To begin with, we construct a mathematical formalism capable of modeling semantics. In essence, semantics naturally emerge as a meaningfully organized ensemble of tokens (Hilpert & Saavedra, 2020; Apidianaki, 2023). Consequently, tokens are considered microscopic (micro) observations with sophisticated patterns in a sentence, whereas semantics represent macroscopic (macro) observations emerging with more predictable behaviors. Inspired by information theory, we formalize the model’s proficiency in semantic understanding, i.e., information emergence, as the difference of the entropy reduction between micro-level and macro-level. In another word, **a better model proficient in deriving semantics from tokens, in comparison to other models, ought to render a higher entropy reduction for a global sequence than for a single token.**

To compute IE in transformer models, as discussed earlier, we need to mathematically measure entropy reduction for both micro and macro levels. Given the auto-regressive nature of the NTP mechanism, at any layer l in transformer, the most micro-level transition can be naturally framed as the probability $p_{h_l^0|h_{l-1}^0}$ for an isolated token h_l^0 , whereas the macro-level transition can be formulated as $p_{h_l^T|h_{l-1}^0, h_{l-1}^1, \dots, h_{l-1}^T}$ across T tokens. We resort to the mutual information between successive transformer layers and adopt a practically effective estimation algorithm motivated by (Belghazi et al., 2018) which is suitable for high-dimensional continuous representations. Therefore, we can measure the IE value for any token at any transformer layer, reflecting the strength of the LLM’s capability in extracting semantics from the historical context.

We devise a suite of comprehensive experiments encompassing two different scenarios. In the first scenario, we curate a group of synthetic datasets under the ICL setting with different context domains. In the second scenario, we collect two wild datasets consisting of real-world natural language questions/answers. Under both scenarios, we experiment with different LMs including GPT-2 (Radford et al., 2019), GEMMA (Team et al., 2024), and OpenLlama (Computer, 2023). In alignment with our hypothesis, we show that IE offers a high-level informativeness through semantic faithfulness and sensitivity - the richer the semantics, the higher the IE. Additionally, a set of experiments conducted across model sizes have indicated the potential association between IE and Emergence. Furthermore, we obtain 3 interesting findings:

1. IE increases token-by-token in natural texts, whereas, in ICL-style texts, IE increases only when a new demonstration appears.
2. There is a strong correlation between specific hallucination phenomenons and a high variance in IE scores.
3. Distinctive patterns in IE have been observed between human-written and LLM-generated texts, revealing IE’s potential in automatically recognising LLM generations.

Overall, the main contributions could be summarized below:

- We introduce IE, a novel, reasonably validated, and task-agnostic metric to deterministically quantify the semantic understanding capability of LLMs.
- We introduce a light-weight implementation method for evaluating IE, which can be applied to extremely large and closed-source LMs like GPT-4 (Achiam et al., 2023).
- Empirical evidence demonstrates that IE can uncover previously unknown and essential patterns in areas such as ICL, Emergence, and hallucination.

2 RELATED WORK

2.1 INFORMATION EMERGENCE AND EMERGENCE

Emergence is defined as a capability that does not exist in smaller models but appears in larger ones (Srivastava et al., 2023; Lu et al., 2023; Yu & Dong, 2022; Liu et al., 2024). Most commonly, as the model size increases, the performance on many tasks rapidly improves. Several studies have posited explanations for Emergence: the qualitative change resulting from quantitative behavior (Wei et al., 2022), the random combination of linguistic skills (Arora & Goyal, 2022), and principles that can be extrapolated from smaller models (Schaeffer et al., 2023).




Figure 1: The analogy of auto-regressive process in NTP to Markov process. Taking the output representation of token2 in Block 0 (h_1^2) as an example, which receives information from input representations of h_0^0 , h_0^1 , and h_0^2 , satisfying $p_{h_{i+1}^t | H_i^{\leq t}} = p_{h_{i+1}^t | h_i^0, h_i^1, h_i^2}$.

IE is a concept defined and validated in Information Theory (Bedau, 1997; 2008). It describes phenomena observable at the macroscopic level but unobservable at the microscopic level. Nevertheless, only limited empirical experiments are conducted to reflect a similar pattern of abrupt improvement between IE and Emergence with increased model size.

2.2 EVALUATION ON LLM CAPABILITIES

The prevalent body of research extensively measures the capabilities of LLMs across various tasks by employing substantial benchmark datasets (Srivastava et al., 2023; Wang et al., 2024; Zhu et al., 2024). Additionally, a significant amount of research focuses on the performance of LLMs concerning specific capabilities such as adaptability to different domains (Afzal et al., 2024), human-like cognition (opinions, attitudes, etc.) (Ma et al., 2024), followed with input instructions (Zeng et al., 2023), text searching capability (Sun et al., 2023), and reasoning ability (Yang et al., 2024). Moreover, numerous studies have also considered the performance of LLMs in collaborative evaluation with humans (Kim et al., 2024; Zheng et al., 2023). In contrast to these studies, our work concentrates on an essential yet abstract ability of LLMs - the ability to extract semantics from tokens.

3 METHOD

3.1 HOW TO MODEL SEMANTIC IN LLMs

In this paper², we identify the transformer block as the fundamental unit. Specifically, we employ $l = 0, 1, \dots, L - 1$ to index transformer blocks within a language model, where L represents the total number of blocks. For instance, GPT-2 XL (1.6B parameters) comprises 12 blocks ($L = 12$), and Gemma-2B totals 18 blocks ($L = 18$). For any transformer block l , given an input sequence of token length T and hidden state dimension D , the input representation is given by $H_l = \{h_l^0, h_l^1, h_l^2, \dots, h_l^{T-1}\}$ and the output representation is $H_{l+1} = \{h_{l+1}^0, h_{l+1}^1, h_{l+1}^2, \dots, h_{l+1}^{T-1}\}$, where $H \in \mathbb{R}^{T \times D}$ and $h^t \in \mathbb{R}^{1 \times D}$. Without loss of generality, we hypothesize that the multi-layer blocks constitutes a Markov process.

Hypothesis 1 (Markov Process Analogy). *The auto-regressive process of NTP mechanism in multi-layer blocks undergoes a Markovian stochastic process following a transition probability of any h_{i+1}^t with $p_{h_{i+1}^t | h_i^0, h_i^1, h_i^2, \dots, h_i^t}$, simply denoted by $p_{h_{i+1}^t | H_i^{\leq t}}$.*

To simplify the analogy for easier understanding, Figure 1 omits normalization layers, MLP layers, and residual structures between transformer blocks, and thus the output of l -th block is directly considered as the input to $l + 1$ -th block (H_{l+1}). However, in all our real implementations, we retain the exact transformer output at every layer, i.e., $h_{l+1} = h_l + \text{attention}(h_l) + \text{MLP}(h_l)$.

Accordingly, we could categorize token variables within each sequence into two distinct categories: **microscopic (micro) variables** and **macroscopic (macro) variables**. A micro variable refers to a

²Please note that this paper addresses decoder-only autoregressive language models, while other language models that do not meet the requirements, such as BERT (Devlin et al., 2018) family, are not considered within the scope.

162 token which is solely influenced by a single token as the input. For instance, h^0 satisfies $p_{h_{l+1}^0|h_l^0}$.
 163 Whereas macro variables aggregate information from all micro variables and thus encompass tokens
 164 which are influenced by all the tokens within the sequence as the input. An example could be h_{l+1}^{T-1}
 165 which satisfies $p_{h_{l+1}^{T-1}|h_l^0, h_l^1, \dots, h_l^{T-1}}$.
 166

167 In summary, the NTP mechanism can be viewed as a behavior that increasingly coarsens from the
 168 most micro to the most macro scale and finally forms meaningful semantics. Hence, the macro level
 169 represents the semantics level and the micro level indicates the token level. Our definition of IE,
 170 hence, represents the phenomenon observable at a semantics level, yet unobservable at a token level
 171 during a dynamic process (instantiated in Example 1).

172 **Example 1.** Given T binary tokens $H_l = \{h_l^0, h_l^1, \dots, h_l^{T-1}\} \in \{0, 1\}^T$ as inputs, for simplicity,
 173 we assume all variables are micro variables: $\forall h_{l+1}^t \in H_{l+1}$ satisfies $p_{h_{l+1}^t|h_l^t}$ (the simplest Markov
 174 process, and in the subsequent part of this Example, we use p to simply denote this transition
 175 probability). The output representations are also binary, i.e., $H_{l+1} \in \{0, 1\}^T$. We assume an
 176 evolution rule which enables the parity of the sum of all output variables equal to the sum of all inputs
 177 with probability γ . If H_l satisfies the uniform distribution, the evolution rule entails the probability of
 178 the output H_{l+1} :

$$179 \quad p(H_{l+1}|H_l) = \begin{cases} \frac{\gamma}{2^{T-1}}, & \text{if } \bigoplus_{t=0}^{T-1} h_{l+1}^t = \bigoplus_{t=0}^{T-1} h_l^t \\ \frac{1-\gamma}{2^{T-1}}, & \text{otherwise} \end{cases} \quad (1)$$

181 where $\bigoplus_{t=0}^{T-1} h_l^t := 1$ if $\sum_{t=0}^{T-1} h_l^t$ is even and $\bigoplus_{t=0}^{T-1} h_l^t := 0$ if odd. For example, if $H_l = \{0, 0, 0\}$, H_{l+1}
 182 can be one of $\{0, 0, 0\}, \{0, 1, 1\}, \{1, 0, 1\}, \{1, 1, 0\}$ with probability γ , leading to $\frac{\gamma}{2^3-1}$ chance for
 183 each candidate above. Each of the remaining value for H_{l+1} has probability $\frac{1-\gamma}{2^3-1}$.
 184

185 With the assumption of micro dependency $p_{h_{l+1}^t|h_l^t}$, we can derive the distribution of a **micro variable**
 186 as $p(h_{l+1}^t=0|h_l^t) = p(h_{l+1}^t=1|h_l^t) = 0.5$. Finally, let h^{ma} be the macro variable with $h^{ma} = \bigoplus_{t=0}^{T-1} h^t$.
 187 Then the distribution of the **macro variable** becomes:

$$188 \quad p(h_{l+1}^{ma}|h_l^{ma}) = \begin{cases} \gamma, & \text{when } h_{l+1}^{ma} = h_l^{ma} \\ 1 - \gamma, & \text{when } h_{l+1}^{ma} \neq h_l^{ma} \end{cases} \quad (2)$$

192 Example 1 elucidates an interesting phenomenon of IE: The macro variable h^{ma} is not induced by
 193 any individual micro variable but a collective of them. As a result, it shows different phenomenon
 194 from any micro variable. Moreover, we note that the same evolution rule applies to both micro and
 195 macro variables. This is known as **Supervenience Hypothesis** (Bedau, 1997; 2008):

196 **Hypothesis 2** (Supervenience). *When the properties of the micro-level mechanisms of a system are*
 197 *fixed, so are the properties of all its macro levels.*

199 Hypothesis 2 explicates that the ‘‘new phenomenon’’ in macro systems does not materialize *ex nihilo*.
 200 The micro level is causally and mechanically complete, and there is no room left for any causal and
 201 mechanical contribution at the macro level. In other words, the semantics that emerged in macro
 202 has irreducible causal and mechanical power in practice but not in principle, just because the token
 203 perspective is too fine-grained to observe such a phenomenon.

204 Given the Supervenience hypothesis and motivated by more endeavors about information theory
 205 (Rosas et al., 2020; Hoel et al., 2013; 2016), we can define IE in LLMs as:

206 **Definition 1** (Information Emergence in LLMs). *For any transformer block l , let h_l^{ma} be the **macro***
 207 *variable (h_l^{ma} satisfies $p_{h_l^{ma}|H_{l \leq T}}$) and let h_l^{mi} be the **micro variable** (h_l^{mi} satisfies $p_{h_{l+1}^{mi}|h_l^{mi}}$),*
 208 *$MI(\cdot, \cdot)$ represents the mutual information, thus the strength of IE in block l can be described as:*

$$209 \quad E(l) = MI(h_{l+1}^{ma}, h_l^{ma}) - \frac{1}{T} \sum_{t=0}^{T-1} MI(h_{l+1}^{mi,t}, h_l^{mi,t}) \quad (3)$$

214 Definition 1 describes how to estimate the IE metric. To illustrate, suppose an input sequence
 215 contains three tokens ‘large language model’, with their representations at the l th block denoted
 as $H_l = \{h_l^0, h_l^1, h_l^2\}$. To compute the mutual information at the micro level, we need to make

sure each micro token is positioned at the beginning of the sequence to avoid the influence from other tokens due to the auto-regressive nature of the NTP mechanism. Specifically, we obtain the first micro variable $\{h_l^{mi.0}\}_{l=0}^{L-1}$ where $h^{mi.0} = h^0$ corresponds to ‘large’ derived by going through the transformer model for the input ‘large language model’. The second micro variable $\{h_l^{mi.1}\}_{l=0}^{L-1}$ corresponds to ‘language’ given by going through the transformer model one more time for the modified input ‘language model’. The third micro variable proceeds in a similar manner by removing the first two tokens in the original input. These modified inputs ensure that each micro variable only depends on itself in the previous block. Meanwhile, the macro variable $\{h_l^{ma}\}_{l=0}^{L-1}$ is given by $h_l^{ma} = h_l^2$ for the original input sequence ‘large language model’. Finally, we have $E(l) = MI(h_{l+1}^{ma}, h_l^{ma}) - \frac{1}{3}(MI(h_{l+1}^{mi.0}, h_l^{mi.0}) + MI(h_{l+1}^{mi.1}, h_l^{mi.1}) + MI(h_{l+1}^{mi.2}, h_l^{mi.2}))$.

Notably, We measure IE by comparing the macro changes with the micro changes across blocks, instead of directly computing $MI(h_l^{ma}, h_l^{mi})$ at the same block, justified by the Supervenience hypothesis which posits that under the same state, micro and macro variables do not exhibit emergence phenomenon due to the completeness of the micro level. We use Example 2 to explain why the mutual information between macro variables and micro variables is different:

Example 2. Continuing from Example 1, for simplicity, we assume $\gamma=1, T=3$. Assuming block l undergoes a uniform distribution, we have $p(h_l^{ma}=0) = p(h_l^{ma}=1) = 0.5$ and $p(h_l^{mi.t}=0) = p(h_l^{mi.t}=1) = 0.5$. Without loss of generality, let’s assume $h_l^{ma} = 0$ and $h_l^{mi.t} = 0$. According to Example 1, we have $p(h_{l+1}^{mi.t}=0|h_l^{mi.t}) = p(h_{l+1}^{mi.t}=1|h_l^{mi.t}) = 0.5$. On the contrary, $p(h_{l+1}^{ma} = 0|h_l^{ma}) = \gamma = 1$. Hence, the IE value becomes $E(l) = 1 - \frac{1}{3}(0 + 0 + 0) = 1$.

Example 2 elucidates that the difference in mutual information stems from a coarse-graining that transitions from individual elements (micro) to an aggregate whole (macro). From an information theory perspective, $E(l) > 0$ indicates that when the function of transformer block l results in a higher reduction of uncertainty (entropy) on the whole sequence (macro variable) compared to the individual tokens (micro variables), there is a higher chance of capturing the collective semantics. Consequently, IE can be briefly understood as **“how confidence with which a language model, based on previous tokens, definitely predicts the next token with a lower entropy in semantics”**.

3.2 HOW TO ESTIMATE IE IN LLMs?

It is not feasible to directly compute the mutual information in Eq. 3 using Kullback-Leibler (KL) divergence, as the input lies in a high-dimensional continuous space. To address that, we resort to an approximation method using mean values proposed in Belghazi et al. (2018):

$$D_{KL}(\mathbb{P}||\mathbb{Q}) = \limsup_{f:\Omega \rightarrow \mathbb{R}} E_{\mathbb{P}}[f] - \log(E_{\mathbb{Q}}[e^f]) \quad (4)$$

where f represents a function that maps \mathbb{Q} to Gibbs distribution by $d\mathbb{G} = \frac{1}{Z} e^f d\mathbb{Q}$, where $Z = E_{\mathbb{Q}}[e^f]$. Naturally, f can be a neural model. Thereby, Equation 4 can be equivalently represented as optimizing the error function \mathcal{L} :

$$\mathcal{L} = \frac{1}{B} \sum_{b=1}^B (f_{\theta}(x^b||y^b)) - \log\left(\frac{1}{B} \sum_{b=1}^B e^{f_{\theta}(x^b||y^{i \neq b})}\right) \quad (5)$$

where θ denotes the parameters for f , $||$ denotes the concatenation operation and B is the batch size. $x, y \in \mathbb{R}^D$ are two inputs for computing the mutual information $MI(x, y)$. $x^b||y^b$ corresponds to sampling from the joint distribution P_{XY} , while $x^b||y^{i \neq b}$ corresponds to sampling from the marginal distribution P_X and P_Y ³. When \mathcal{L} converges to the minimum $\hat{\mathcal{L}}$, we can obtain the final estimated mutual information as $MI(x, y) = -\log_2^e * \hat{\mathcal{L}}$. (More details and proofs are shown in Belghazi et al. (2018).)

To get the IE value $E(l)$ in Eq. 3, we compute $MI(h_{l+1}^{ma}, h_l^{ma})$ by replacing x^b and y^b in Eq. 5 with $h_{l+1,s}^{ma}$ and $h_{l,s}^{ma}$ obtained by applying the LM to the same input sequence s , whereas replacing $y^{i \neq b}$ with $h_{l,s'}^{ma}$ using a different sequence $s' \neq s$. Similar operations apply when computing $MI(h_{l+1}^{mi.t}, h_l^{mi.t})$. (Refer to Appendix A for a complete algorithm for estimator.)

³In our implementation, the batch size was increased to encompass the entirety of the sample set to ensure the rationality of P_{XY}, P_X and P_Y .

3.3 IMPLEMENTATION

Our algorithm requires that the number of samples is sufficiently large (over 300k) to provide a good estimate of the mutual information. Meanwhile, the length of each sequence within a dataset should be kept the same to facilitate position-wise observation and meaningful computation. Due to resource constraints, our comparative analysis is limited to GPT2-large (812M), GPT2-XL (1.61B), GEMMA (2.51B), and OpenLlama (3B) models. Fortunately, this parameter range is sufficient to observe variations and regularities of IE. For those LLMs with extremely large size or closed resource (e.g., GPT-4, Claude3, etc.), we design another efficient strategy that enables their IE evaluations as shown in Section 5.3. All computational experiments can be conducted on **one** NVIDIA GeForce RTX 3090 GPU⁴. The estimator f in Eq. 4 is a model of a 10-layer neural network comprising linear layers and leaky ReLU activation functions, where each linear layer’s output dimension was half of its input dimension. We set the batch size to 300,000 to ensure the stability of the sampling distribution, thereby guaranteeing robust results. Additionally, the learning rate was initially set at 1e-4 and was polynomially decayed to 1e-8 within 10k epochs. We examine the IE value of LLMs under two distinct settings: **ICL** with few-shot examples and **natural sentences** without demonstrations.

3.3.1 ICL SCENARIO

Since existing datasets (e.g., SST-2 (Socher et al., 2013), AGNews (Zhang et al., 2015), and EmoC (Chatterjee et al., 2019)) do not meet the requirement of the same sequence length, we synthesized a set of simple few-shot samples having token length and positions aligned across different sequences. We curate three different datasets encompassing three different domains, each containing sequences of few-shot single-token entities with commas:

- **Country:** We select 25 countries from the Vocabulary as *entities*, each represented by 1 token (e.g., ‘Canada’, ‘Russia’). Each shot consists of one *entity* followed by a *comma*, totaling 2 tokens. We constructed $25 * 24 * 23 * 22 = 303,600^5$ input sequences, each comprising 8 tokens (4 different shots), such as “France, Mexico, Egypt, Russia,”.
- **Animal:** Similarly, we select 16 animals as *entities*, and construct $16 * 15 * 14 * 13 * 12 = 524160$ input sequences comprising 10 tokens (5 different shots), such as “Fox, Pig, Penguin, Rabbit, Cock,”.
- **Color:** We select 15 colors as *entities*, and construct 360360 samples comprising 10 tokens (5 different shots), such as “red, orange, yellow, green, blue,”.

In the experiment, we observe that each *entity*, treated as a micro variable (i.e., the first token), produces similar mutual information across different positions. Consequently, in this section, we only use the *entity* in the first position to compute the mutual information to approximate the mean of all micro variables. Moreover, $E(l)$ also acts analogously in each block, so we utilize the mean of $\{E(l)\}_{l=0}^{L-1}$ to show the IE. By varying the position t of the macro variable (last token in a sequence), we use the following equation to compute the IE at length t :

$$\hat{E}(t) = \frac{1}{L} \sum_{l=0}^{L-1} E(l)_{ma=t-1} = \frac{1}{L} \sum_{l=0}^{L-1} (MI(h_{l+1}^{t-1}, h_l^{t-1}) - MI(h_{l+1}^0, h_l^0)) \quad (6)$$

3.3.2 NATURAL SENTENCE SCENARIO

We randomly select 300,000 natural sequences, each consisting of 8 tokens, from OpenOrca (Lian et al., 2023) and OpenHermes (Teknum, 2023), respectively. OpenOrca and OpenHermes are both large-scale, multi-domain QA datasets. These sequences were selected to ensure that the first token in each sequence is the actual beginning of a sentence.

In our experiments, we observe potential discrepancies in mutual information for individual tokens at different positions within a sentence (i.e., micro mechanisms in different positions are not consistent).

⁴For easy deployment, we split the representations of LLMs into several 768-dimension segments and the final representation is the mean of these segments.

⁵The number of shots is decided to ensure the number of combinations in each category is over 300000.

These discrepancies are detailed in Appendix C). Consequently, for the mutual information of micro-level variables, we adhere to Equation 3 which averages the micro MI at each position:

$$\hat{E}(t) = \frac{1}{L} \sum_{l=0}^{L-1} E(l)_{ma=t-1} = \frac{1}{L} \sum_{l=0}^{L-1} (MI(h_{l+1}^{t-1}, h_l^{t-1}) - \frac{1}{t} \sum_{m=0}^{t-1} MI(h_{l+1}^m, h_l^m)) \quad (7)$$

4 INFORMATIVENESS

4.1 SEMANTIC FAITHFULNESS

We have designed several experiments to substantiate the **semantic faithfulness** of IE. We observed the change in IE and other popular metrics (Exact Match, Accuracy, and model loss) with the increase in the number of tokens, using the samples from the OpenOrca dataset. Figure 2 demonstrates the change in their values, with increasement > 0 representing the value increasing from that in the previous token. Only IE consistently exhibits an upward trend (i.e., > 0), which aligns with the intuition: what a sentence intends to convey is increasingly deterministic along with the increasing number of tokens. Moreover, the low variance (reflected as the shaded area in Figure 2) in IE values exhibits commendable stability compared to other metrics.




Figure 2: The increasement of IE, EM, Accuracy, and model loss for GPT2-XL in comparison to the previous token: increasement = $(value(t) - value(t - 1))/value(t)$, where $value(t)$ represents the value at token t . Therefore, a positive increasement (> 0) indicates an increase in the metric value, and a decrease vice versa.

4.2 SEMANTIC SENSITIVITY

Subsequently, we aim to examine the **semantic sensitivity** of IE, particularly its ability to reflect differences when minor perturbations are introduced into the semantics. Consequently, we conducted a series of ablation studies to modulate certain factors (such as dataset size, attributes, tasks, and format) individually. We treat the performance of GPT2-XL on the “country” dataset as a baseline. Appendix D details the variations when different factors are changed. It was observed that IE increases with the model’s size. This corroborates the rationale that a model with a larger size generally has better capability to determine semantics. In addition, our study also identifies variations in IE against different tasks and prompts, which also resonates with findings from prior research (Lu et al., 2023; Yu & Dong, 2022; Liu et al., 2024).

4.3 CONNECTION TO EMERGENCE

Moreover, We demonstrate that IE manifests a steep ascend within the parameter range of 10^8 to 10^{10} across 8 arithmetic tasks, which is detailed in Appendix B). Given the confines of computational resources, we were able to select 8 models within the parameter range of GPT2 ($1 * 10^8$), GPT2-large ($7 * 10^8$), GPT2-XL ($1 * 10^9$), Gemma ($2 * 10^9$), OpenLlama ($3 * 10^9$), GPT-J ($6 * 10^9$), Gemma ($7 * 10^9$), and GPT-NeoX ($2 * 10^{10}$). In light of the existing evaluation work, Big-Bench (Srivastava et al., 2023), we discovered the emergent phenomena within the arithmetic tasks emerge within the parameters of $[10^8, 10^{10}]$. Consequently, the association between the performance and IE values of 8 arithmetic tasks was investigated, as shown in Figure 3. For model performance, we directly adopt the default settings of the Big-Bench benchmark. As for IE, we took the average of the IE values of the initial five output tokens to be the final result.

A marked enhancement in task performance occurs once effective parameters reach 10^{10} , thereby showcasing an emergent phenomenon. The average IE experiences a substantial surge within the same parameter range (10^9 to 10^{10}). As a pioneering work proposing a quantitative metric to reflect the level of semantics deterministically, we believe our method could also greatly benefit further research on Emergence.

378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431

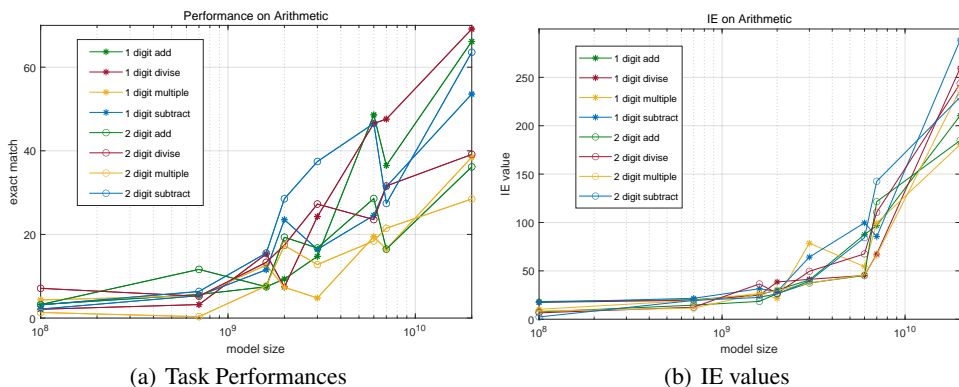


Figure 3: IE and Model Performance with model size increasing in Arithmetic

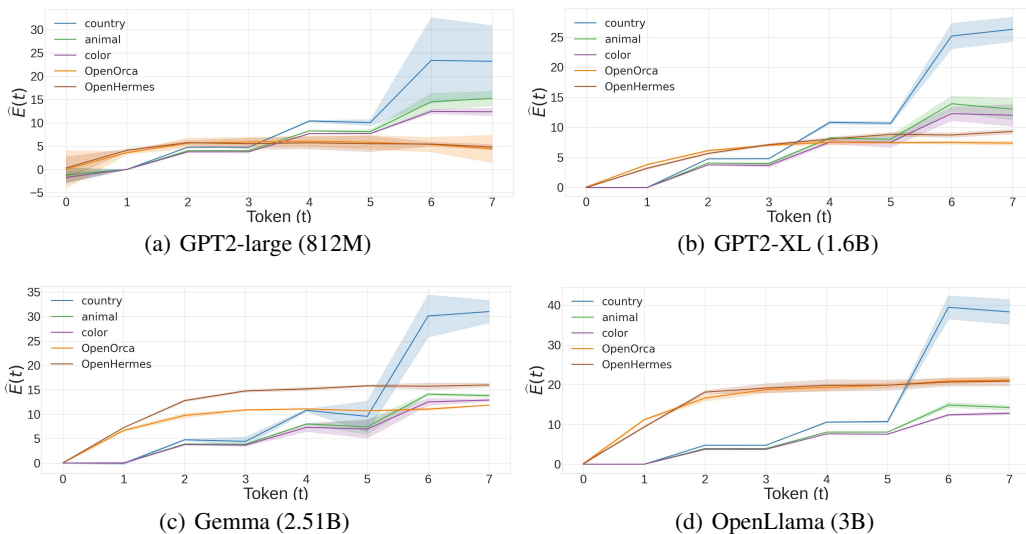


Figure 4: $\hat{E}(t)$ on ICL and natural scenarios with mean and variance.

5 FINDINGS

5.1 IE INCREASES ONLY WHEN A NEW DEMONSTRATION APPEARS

Figure 4 illustrates that IE naturally becomes higher with increased tokens. However, there is a strikingly different trend between ICL and natural scenarios (containing natural sentences). In a natural scenario, IE increases with each successive token and achieves a rapid convergence (around the 6th token), whereas, under the ICL scenario, IE only increases when a new demonstration emerges (at positions of the 2nd token, 4th token, 6th token)⁶, but with a higher upper bound and requiring more tokens to reach to the highest value.

We subsequently investigate how many demonstrations are needed before IE ceases to increase. Table 4 indicates that the three ICL categories under study tend to saturate at the 7th demonstration (though this does not suggest a generic ICL phenomenon). Moreover, we test whether increasing the number of tokens within each demonstration would maintain this ”stepwise elevating” pattern. Figure 6 shows that the IE scores within each demonstration does not change when the length of each demonstration is increased to 5 tokens, 6 tokens, and 7 tokens.

⁶The complete record of every token’s mutual information is detailed in Table 5, showing the example of GPT2-XL on the Animal category.

Hence, we can interpret ICL’s role in enhancing semantic determinability: ICL bolsters semantic determinability via demonstrations, where increasing the number of demonstrations can increase the ability of capturing semantics beyond natural text, but eventually saturates after a certain quantity.

Concurrently, disparate performances observed across the two datasets and three model families suggest that the domain of the training data and preprocessing methodologies are likely critical factors, as further supported by the evaluations of individual tokens at different sentence positions elaborated in Appendix C. For instance, GPT2-large and GPT2-XL, which share the same dataset (40GB WebText), preprocessing methods, and model architecture, exhibit a common characteristic: the mutual information of the first token is consistently lower than the average mutual information of micro-level variables.

5.2 HIGHER IE WITH LARGE STANDARD DEVIATION CORRESPONDS TO CERTAIN HALLUCINATION BEHAVIORS

Figure 4 indicates that IE becomes unstable as the number of demonstrations increases. To further study this observation, we explicitly report the IE and standard deviation (s.d.) in Table 1 and compute the accuracy of the generations⁷ spanning over different numbers of shots. As can be seen from rows 1-9 of Table 1, as IE ceases to grow and the s.d. reaches the peak, the LM displays a higher probability of generating inaccurate responses. From a closer look, we discover that oftentimes, the LM fails to generate new entities due to “error repetition” (explanations and some examples can be found in Appendix E.3). This is aligned with existing study (Zhang et al., 2023) related to hallucination. Specifically, LLMs struggle to correct themselves after generating an erroneous output, consequently resulting in stagnation and fluctuation in IE value.

However, this should not be confused with the power of ICL in exploiting more complex patterns effectively with more “shots” as the input. Different from the above observation, an increasing number of shots tend to bring higher accuracies under more complex scenarios. As shown in rows 10-15 of Table 1, we design four challenging tasks: ‘Asia’ and ‘Europe’ only provide countries in Asia and Europe, respectively, as input demonstrations; ‘Size’ contains animals arranged by size from smaller to larger; ‘Alphabet’ sorts the entities alphabetically based on the first letter. The accuracy results indicate that LLMs require more shots to capture complex patterns compared to simple patterns. Thus, it prompts us to conjecture if the stagnation and fluctuations in the IE are associated with another hallucination: with excessive shots, LLMs may perceive more complex patterns beyond the surface (or even actual) appearance. In short, the correlation between IE s.d. and hallucination would offer novel insights into the future development of hallucination detection and mitigation.

IE value by each shot							
Statistics	shot1	shot2	shot3	shot4	shot5	shot6	shot7
value	4.013	8.34	12.95	26.81	61.59	82.49	71.52
SD	<0.01	0.59	0.84	2.61	6.59	7.22	7.05
Accuracy of LLMs outputs given shots (%)							
dataset	shot1	shot2	shot3	shot4	shot5	shot6	shot7
country	0	54.15	74.29	88.47	46.21	21.59	22.68
animal	0	44.51	69.43	76.19	64.19	36.14	33.54
color	0	37.49	66.51	72.18	73.16	46.95	38.49
Accuracy in 4 complex pattern given shots (%)							
pattern	shot1	shot2	shot3	shot4	shot5	shot6	shot7
Asia	0.35	3.27	4.26	15.29	34.72	84.53	79.16
Europe	3.75	8.29	11.16	24.68	49.36	89.38	89.51
Size	4.59	2.94	6.43	7.29	7.16	26.46	34.19
Alphabet	0.11	1.26	1.47	39.16	69.17	54.91	18.67

Table 1: The relationship between the accuracy of GPT2-XL outputs and IE by each shot in 3 categories.

5.3 TEXTS GENERATED FROM LLMs AND HUMANS EXHIBIT DIFFERENT IE VALUES

We seek to measure the differences in text generated by larger language models compared to human texts, as well as the variations among these LLMs themselves. Specifically, we use questions from OpenHermes as inputs and collect responses by invoking the APIs of GPT-4, Claude3-opus, Claude3-sonnet, and Llama3. These responses were subjected to the same data processing methods described in Section 3.3.2. To evaluate these extremely large and closed-source language models, we implement a 3-step strategy:

⁷We randomly sample a total of $N = 1000$ samples and regard the generation to be correct if the generated entity belongs to the corresponding domain of the dataset (country, animal or color) and is not repetitive.

Text+Estimator	token0	token1	token2	token3	token4	token5	token6	token7	token8
Human+GPT2-XL	10.9	16.9	18.6	19.5	19.5	19.7	19.6	19.5	19.4
Human+GEMMA	9.5	16.8	22.4	24.3	24.0	25.3	24.6	25.0	25.9
GPT4+GPT2-XL	11.3	18.8	23.5	27.2	34.5	37.2	39.2	39.5	39.2
GPT4+GEMMA	12.1	20.5	25.1	31.6	36.3	39.9	40.4	39.5	40.6
Claude3-opus+GPT2-XL	12.6	21.8	26.6	29.5	36.8	39.8	42.6	45.2	45.3
Claude3-sonnet+GPT2-XL	11.4	17.4	24.8	28.5	32.5	36.5	36.1	36.2	36.2
Llama3 (70B)+GPT2-XL	11.2	18.1	23.6	24.5	28.5	32.6	36.5	36.8	36.6

Table 2: IE in texts generated from human and popular LLMs. “text” refers to the party that generates the text. “Estimator” refers to the LM used to transform the text into representations and estimates the IE value using f described in Section 3.2.

Step 1: Collect the answers from these LLMs (or humans) via the questions from the OpenHermes.

Step 2: Following the data processing in Section 3.3.2, we format these answers into input sequences of 8 tokens and obtained their representations using smaller LMs (e.g., GPT-2, GEMMA).

Step 3: These representations were then processed through an estimator to calculate the mutual information introduced in Section 3.2, thereby determining the IE values of these answers via Equation 7.

Table 2 illustrates an interesting phenomenon: LLM-generated texts exhibited substantially greater IE value than human texts. This observation is intuitive—given that LLMs aim to generate tokens with the highest probability, naturally resulting in greater entropy reduction.

Another observation is that the text generated by different LLMs (GPT-4, Claude3, and Llama3) displays variations in IE values. Significant differences are observed not only in the maximum strength of the IE but also in the patterns of growth. Without actually computing the transformer representations of the target LLMs, these findings open a promising path towards estimating the semantic capturing capability from extremely large and closed-source LMs without expensive computational costs.

6 LIMITATIONS

Position-wise Token: Given that mutual information intrinsically demands the distribution of two tokens to be valid, we require every token’s position to hold a specific meaning, such as representing the beginning or end of a sentence, the subject, predicate, and so forth. Hence, applying our estimator directly to existing tasks may result in a lack of interpretability as the token lengths and positions in the samples vary significantly.

Sample Amount: To ensure the accuracy of joint and marginal distributions of high-dimensional continuous representations, a tremendous number of samples is essential. We are attempting more mechanistic alternative methods, hoping to reduce sample size requirements in the future.

More Models and Tokens: It is evident that our experiments lack larger-sized models and analysis of long-length texts, especially for emergence and hallucination analysis. Given more computational resources, we would continue to expand these experiments.

7 CONCLUSION

In this paper, we mathematically model the entropy of tokens and propose a quantitative metric, Information Emergence (IE), representing the LLM’s ability to obtain semantics from tokens. Under the proposed low-resource estimator, we corroborate that IE possesses semantic faithfulness and sensitivity not found in other metrics. Under the settings of ICL and natural sentences, We conducted extensive experiments explaining why ICL provides clearer semantics than natural text, as well as the intrinsic relationship between IE and hallucinations. Simultaneously, we discovered that IE can be utilized to distinguish whether the source of the text originates from humans or LLMs, particularly a simple and feasible strategy for those of significantly large LMs.

REFERENCES

- 540
541
542 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman,
543 Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report.
544 *arXiv preprint arXiv:2303.08774*, 2023.
- 545 Anum Afzal, Ribin Chalumattu, Florian Matthes, and Laura Mascarell. Adapteval: Evaluating large
546 language models on domain adaptation for text summarization, 2024. URL <https://arxiv.org/abs/2407.11591>.
- 547
548 Marianna Apidianaki. From word types to tokens and back: A survey of approaches to word meaning
549 representation and interpretation. *Computational Linguistics*, 49(2):465–523, 2023.
- 550
551 Sanjeev Arora and Anirudh Goyal. A theory for emergence of complex skills in language models.
552 *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. Survey Certification.
- 553
554 Mark A Bedau. Weak emergence. *Philosophical perspectives*, 11:375–399, 1997.
- 555
556 Mark A Bedau. Downward causation and autonomy in weak emergence. 2008.
- 557 Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeshwar, Sherjil Ozair, Yoshua Bengio, Aaron
558 Courville, and Devon Hjelm. Mutual information neural estimation. In *International conference*
559 *on machine learning*, pp. 531–540. PMLR, 2018.
- 560 Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan
561 Yi, Cunxiang Wang, Yidong Wang, et al. A survey on evaluation of large language models. *ACM*
562 *Transactions on Intelligent Systems and Technology*, 15(3):1–45, 2024.
- 563
564 Ankush Chatterjee, Kedhar Nath Narahari, Meghana Joshi, and Puneet Agrawal. Semeval-2019 task
565 3: Emocontext contextual emotion detection in text. In *Proceedings of the 13th international*
566 *workshop on semantic evaluation*, pp. 39–48, 2019.
- 567 Anthony Chen. *Factuality and Large Language Models: Evaluation, Characterization, and Corre-*
568 *ction*. University of California, Irvine, 2023.
- 569
570 Together Computer. Redpajama-data: An open source recipe to reproduce llama training dataset,
571 2023. URL <https://github.com/togethercomputer/RedPajama-Data>.
- 572
573 Kees van Deemter. The pitfalls of defining hallucination. *Computational Linguistics*, 50(2):807–816,
574 2024.
- 575
576 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep
bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- 577
578 Martin Hilpert and David Correia Saavedra. Using token-based semantic vector spaces for corpus-
579 linguistic analyses: From practical applications to tests of theoretical claims. *Corpus Linguistics*
580 *and Linguistic Theory*, 16(2):393–424, 2020.
- 581
582 Erik P Hoel, Larissa Albantakis, and Giulio Tononi. Quantifying causal emergence shows that macro
can beat micro. *Proceedings of the National Academy of Sciences*, 110(49):19790–19795, 2013.
- 583
584 Erik P Hoel, Larissa Albantakis, William Marshall, and Giulio Tononi. Can the macro beat the micro?
585 integrated information across spatiotemporal scales. *Neuroscience of Consciousness*, 2016(1):
586 niw012, 2016.
- 587
588 Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang,
589 Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM*
Computing Surveys, 55(12):1–38, 2023.
- 590
591 Tae Soo Kim, Yoonjoo Lee, Jamin Shin, Young-Ho Kim, and Juho Kim. Evallm: Interactive
592 evaluation of large language model prompts on user-defined criteria. In *Proceedings of the CHI*
593 *Conference on Human Factors in Computing Systems*, CHI ’24, New York, NY, USA, 2024.
Association for Computing Machinery. ISBN 9798400703300. doi: 10.1145/3613904.3642216.
URL <https://doi.org/10.1145/3613904.3642216>.

- 594 Jannik Kossen, Yarin Gal, and Tom Rainforth. In-context learning learns label relationships but is
595 not conventional learning. In *The Twelfth International Conference on Learning Representations*,
596 2023.
- 597 Wing Lian, Bley Goodson, Eugene Pentland, Austin Cook, Chanvichet Vong, and "Teknium".
598 Openorca: An open dataset of gpt augmented flan reasoning traces. <https://huggingface.co/Open-Orca/OpenOrca>, 2023.
- 601 Zeyu Liu, Yitong Liu, Zehao Zhang, Lei Di, Feng Wei, and Yin Wang. Method for extracting
602 power emergency plan information based on llm prompt learning. In *International Conference on*
603 *Frontiers of Applied Optics and Computer Engineering (AOCE 2024)*, volume 13080, pp. 106–111.
604 SPIE, 2024.
- 605 Sheng Lu, Irina Bigoulaeva, Rachneet Sachdeva, Harish Tayyar Madabushi, and Iryna Gurevych.
606 Are emergent abilities in large language models just in-context learning? *arXiv preprint*
607 *arXiv:2309.01809*, 2023.
- 609 Bolei Ma, Xinpeng Wang, Tiancheng Hu, Anna-Carolina Haensch, Michael A. Hedderich, Barbara
610 Plank, and Frauke Kreuter. The potential and challenges of evaluating attitudes, opinions, and
611 values in large language models, 2024. URL <https://arxiv.org/abs/2406.11096>.
- 612 Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hannaneh Hajishirzi. MetaICL: Learning to learn
613 in context. In Marine Carpuat, Marie-Catherine de Marneffe, and Ivan Vladimir Meza Ruiz
614 (eds.), *Proceedings of the 2022 Conference of the North American Chapter of the Association for*
615 *Computational Linguistics: Human Language Technologies*, pp. 2791–2809, Seattle, United States,
616 July 2022a. Association for Computational Linguistics. doi: 10.18653/v1/2022.naacl-main.201.
617 URL <https://aclanthology.org/2022.naacl-main.201>.
- 618 Sewon Min, Xinxin Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke
619 Zettlemoyer. Rethinking the role of demonstrations: What makes in-context learning work? In
620 *EMNLP*, 2022b.
- 622 Shervin Minaee, Tomas Mikolov, Narjes Nikzad, Meysam Chenaghlu, Richard Socher, Xavier
623 Amatriain, and Jianfeng Gao. Large language models: A survey. *arXiv preprint arXiv:2402.06196*,
624 2024.
- 625 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language
626 models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- 628 Vipula Rawte, Swagata Chakraborty, Agnibh Pathak, Anubhav Sarkar, SM Towhidul Islam Tonmoy,
629 Aman Chadha, Amit Sheth, and Amitava Das. The troubling emergence of hallucination in
630 large language models—an extensive definition, quantification, and prescriptive remediations. In
631 *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp.
632 2541–2573, 2023.
- 633 Fernando E Rosas, Pedro AM Mediano, Henrik J Jensen, Anil K Seth, Adam B Barrett, Robin L
634 Carhart-Harris, and Daniel Bor. Reconciling emergences: An information-theoretic approach to
635 identify causal emergence in multivariate data. *PLoS computational biology*, 16(12):e1008289,
636 2020.
- 637 Rylan Schaeffer, Brando Miranda, and Sanmi Koyejo. Are emergent abilities of large language
638 models a mirage? In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
639 URL <https://openreview.net/forum?id=ITw9edRD1D>.
- 641 Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and
642 Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank.
643 In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pp.
644 1631–1642, 2013.
- 645 Aarohi Srivastava, Abhinav Rastogi, and et al Abhishek Rao. Beyond the imitation game: Quantifying
646 and extrapolating the capabilities of language models. *Transactions on Machine Learning Research*,
647 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=uyTL5Bvosj>.

- 648 Weiwei Sun, Lingyong Yan, Xinyu Ma, Shuaiqiang Wang, Pengjie Ren, Zhumin Chen, Dawei Yin,
649 and Zhaochun Ren. Is chatgpt good at search? investigating large language models as re-ranking
650 agents. *arXiv preprint arXiv:2304.09542*, 2023.
- 651
- 652 Sivaramakrishnan Swaminathan, Antoine Dedieu, Rajkumar Vasudeva Raju, Murray Shanahan,
653 Miguel Lazaro-Gredilla, and Dileep George. Schema-learning and rebinding as mechanisms of
654 in-context learning and emergence. *Advances in Neural Information Processing Systems*, 36, 2024.
- 655 Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak,
656 Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. Gemma: Open models
657 based on gemini research and technology. *arXiv preprint arXiv:2403.08295*, 2024.
- 658
- 659 Teknium. Openhermes 2.5: An open dataset of synthetic data for generalist llm assistants, 2023.
660 URL <https://huggingface.co/datasets/teknium/OpenHermes-2.5>.
- 661 Yidong Wang, Zhuohao Yu, Wenjin Yao, Zhengran Zeng, Linyi Yang, Cunxiang Wang, Hao Chen,
662 Chaoya Jiang, Rui Xie, Jindong Wang, Xing Xie, Wei Ye, Shikun Zhang, and Yue Zhang. PandaLM:
663 An automatic evaluation benchmark for LLM instruction tuning optimization. In *The Twelfth
664 International Conference on Learning Representations*, 2024. URL [https://openreview.
665 net/forum?id=5Nn2BLV7SB](https://openreview.net/forum?id=5Nn2BLV7SB).
- 666 Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani
667 Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto,
668 Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. Emergent abilities of large lan-
669 guage models. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. URL
670 <https://openreview.net/forum?id=yzkSU5zdwD>. Survey Certification.
- 671
- 672 Noam Wies, Yoav Levine, and Amnon Shashua. The learnability of in-context learning. *Advances in
673 Neural Information Processing Systems*, 36, 2024.
- 674 Menglin Yang, Aosong Feng, Bo Xiong, Jiahong Liu, Irwin King, and Rex Ying. Enhancing llm
675 complex reasoning capability through hyperbolic geometry. In *ICML 2024 Workshop on LLMs
676 and Cognition*, 2024.
- 677
- 678 Jiacheng Ye, Zhiyong Wu, Jiangtao Feng, Tao Yu, and Lingpeng Kong. Compositional exemplars for
679 in-context learning. In *International Conference on Machine Learning*, pp. 39818–39833. PMLR,
680 2023.
- 681 Zhibin Yu and Yanping Dong. The emergence of a complex language skill: Evidence from the
682 self-organization of interpreting competence in interpreting students. *Bilingualism: Language and
683 Cognition*, 25(2):269–282, 2022. doi: 10.1017/S1366728921000870.
- 684 Zhiyuan Zeng, Jiatong Yu, Tianyu Gao, Yu Meng, Tanya Goyal, and Danqi Chen. Evaluating large
685 language models at evaluating instruction following. In *The Twelfth International Conference on
686 Learning Representations*, 2023.
- 687
- 688 Muru Zhang, Ofir Press, William Merrill, Alisa Liu, and Noah A. Smith. How language model
689 hallucinations can snowball, 2023.
- 690
- 691 Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text
692 classification. *Advances in neural information processing systems*, 28, 2015.
- 693
- 694 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang,
695 Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica.
696 Judging LLM-as-a-judge with MT-bench and chatbot arena. In *Thirty-seventh Conference on
697 Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. URL [https://
698 openreview.net/forum?id=uccHPGDlao](https://openreview.net/forum?id=uccHPGDlao).
- 699
- 700 Lianghui Zhu, Xinggang Wang, and Xinlong Wang. JudgeLM : Fine-tuned large language models
701 are scalable judges, 2024. URL <https://openreview.net/forum?id=87YOFayjcG>.

A ALGORITHM FOR ESTIMATING MUTUAL INFORMATION

Algorithm 1 Estimating Mutual Information

Require: : A set of input tokens $U \in \mathbb{R}^{S*T}$, where S denotes the total number of samples and T represents the number of token in each sample. A LLM f_τ with L layer of blocks and hidden state dimension D . A estimator f_θ .

Ensure: : Mutual Information $M \in \mathbb{R}^{L*T}$.

procedure 1 Extracting Representation $H \in \mathbb{R}^{S*L*T*D}$ from LLM
Initialization: $H = \emptyset$
for each sample s in S **do**
 $H \leftarrow H + f_\tau(U_s)$
end for

procedure 2 Estimating Mutual Information M
Initialization: $M = \emptyset, l = 0, t = 0,$
while $l < L$ and $t < T$ **do**
 $I_x \leftarrow H_{l,t}(H_{l,t} \in \mathbb{R}^{S*D})$
 $I_y \leftarrow H_{l+1,t}(H_{l+1,t} \in \mathbb{R}^{S*D})$
 Shuffle $H_{l+1,t}$ in the dimension S
 $I_{y^*} \leftarrow H_{l+1,t}(H_{l+1,t} \in \mathbb{R}^{S*D})$
 \leftarrow I_x || I_y
 \leftarrow I_x || I_{y^*}
 Initialization: $M_{tmp} = 0$
 for Epoch $i < 10k$ **do**
 \leftarrow f_\theta(input1)
 \leftarrow f_\theta(input2)
 $\mathcal{L} = \frac{1}{S} \sum_{s=1}^S (output1) - \log(\frac{1}{S} \sum_{s=1}^S (output2))$
 \mathcal{L} backpropagation
 if $M_{tmp} == 0$ **then**
 $M_{tmp} \leftarrow -\log_2^e \mathcal{L}$
 else if $M_{tmp} \neq 0$ **then**
 if $M_{tmp} < -\log_2^e \mathcal{L}$ **then**
 $M_{tmp} \leftarrow -\log_2^e \mathcal{L}$
 end if
 end if
 end for
 $l \leftarrow l + 1, t \leftarrow t + 1$
 $M_{l,t} \leftarrow M_{tmp}$
end while
return M

Algorithm 1 is employed to elucidate the entire process of estimating mutual information. Simplified, the method involves two primary steps: Step 1 involves extracting representative samples from a LLM, and Step 2 entails estimating the mutual information between these representation samples. We denote the time required to estimate a pair of representations ($H_{l,t}$ and $H_{l+1,t}$) as α . Consequently, the time complexity for estimating representations from an LLM for a sequence $S_T = token1, token2, ;tokenT - 1$ across L block layers is denoted as $O(LT\alpha)$.

In practical implementations, α approximately costs 40 minutes on one 3090 GPU, whereas significant improvements on a 4090 GPU reduce this time to about 20 minutes.

B CASES IN ARITHMETIC TASKS

We have selected a total of 8 arithmetic tasks, as illustrated in the caption of Figure 3. For these tasks, we employed the 2-shots as the prompt templates for the ICI method. We randomly matched different shots for each sample. A representative example from each task is selected and presented as follows:

1 digit addition:

756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

OpenOrca

	token0	token1	token2	token-3	token-2	token-1
with previous token	8.38	13.42	15.56		17.11	17.16	17.26
wo. previous token	8.38	10.75	10.63		10.69	10.64	12.85

OpenHermes

	token0	token1	token2	token-3	token-2	token-1
with previous token	7.91	12.37	15.03		18.37	18.43	19.22
wo. previous token	7.91	10.78	10.91		10.88	10.95	12.62

(a) GPT2-XL

OpenOrca

	token0	token1	token2	token-3	token-2	token-1
with previous token	10.53	17.27	20.17		21.27	21.66	22.41
wo. previous token	10.53	10.65	10.49		10.58	10.48	10.62

OpenHermes

	token0	token1	token2	token-3	token-2	token-1
with previous token	9.59	16.80	22.28		25.31	25.33	25.54
wo. previous token	9.59	9.66	9.62		9.48	9.67	9.61

(b) GEMMA

Figure 5: Mutual information of each token position in two datasets, taking GPT2-XL and GENNA as examples.

measure	t0	t1	t2	t3	t4	t5	t6	t7
baseline	4.69	4.69	9.46	9.37	15.32	15.02	28.44	29.47
model1	4.64	4.69	9.44	9.45	15.09	14.67	27.59	28.67
model2	4.64	4.68	9.44	9.28	15.27	14.66	44.08	36.37
model3	4.69	4.68	9.47	9.45	15.29	15.54	52.28	85.61
token1	2.82	2.83	6.88	6.85	11.08	10.95	16.83	16.09
token2	3.60	3.60	7.36	7.29	11.15	11.08	15.86	14.96
candidate	3.26	3.26	7.22	7.22	11.44	11.35	17.15	17.33
fusion1	3.84	3.84	7.63	7.64	11.66	11.49	17.45	16.28
fusion2	3.45	3.45	7.26	7.05	11.05	11.06	16.45	16.59
space	4.69	4.69	9.46	9.37	15.32	15.02	22.44	5.19
prefix	4.69	4.69	9.46	9.46	15.32	15.34	37.85	41.05

Table 3: “Ablation Study” of how IE value changes with different measures adopted. t0-t7 represent 1st - 8th token.

“What is 1 plus 0? A: 1, What is 4 plus 4? A: 8, What is 2 plus 7? A:”

1 digit division:

“What is 6 divided by 1? A: 6, What is 8 divided by 4? A: 2, What is 3 divided by 3? A:”

1 digit multiplication:

“What is 1 times 8? A: 8, What is 5 times 0? A: 0, What is 6 times 7? A:”

1 digit subtraction:

“What is 5 minus 2? A: 3, What is 7 minus 6? A: 1, What is 9 minus 0? A:”

2 digit addition:

“What is 53 plus 97? A: 150, What is 89 plus 25? A: 114, What is 75 plus 63? A:”

2 digit division:

“What is 72 divided by 9? A: 8, What is 81 divided by 27? A: 3, What is 18 divided by 3? A:”

2 digit multiplication:

“What is 95 times 55? A: 5225, What is 92 times 88? A: 8096, What is 43 times 42? A:”

2 digit subtraction:

“What is 25 minus 14? A: 11, What is 55 minus 36? A: 19, What is 80 minus 38? A:”

C MUTUAL INFORMATION IN THE NATURAL SCENARIO

We observed variations in the IE statistics of tokens at different positions within a sentence. Consequently, we systematically evaluated tokens at various positions within a sentence, as illustrated in Figure 5. Specifically, token0, token1, and token2 were derived from the same sample set A from OpenOrca, while token-3, token-2, and token-1 were taken from another sample set B from OpenOrca. Sample set A ensured that token0 was the initial token of the sentence, while set B ensured that token-1 was the last token of the sentence. This allowed us to measure differences in IE statistics for tokens at the beginning, middle, and end of sentences across variable sentence lengths.

Figure 5 presents an interesting phenomenon: taking GPT2-XL and GEMMA as examples, GPT2-XL exhibits distinct responses to tokens at different sentence positions—IE values increase at the beginning, stabilize in the middle, and rise again at the end. GEMMA, on the other hand, does not display such positional sensitivity. We hypothesize that this may be related to the different preprocessing methods used in the training data.

Model	categories	shot3	shot4	shot5	shot6	shot7
GPT2-large	country	↑5.67	↑12.95	↑2.75	↑0.33	↑1.04
	animal	↑4.24	↑6.06	↑9.52	↑0.39	↓1.44
	color	↑4.88	↑4.82	↑6.39	↑1.24	↑0.22
GPT2-XL	country	↑5.86	↑13.12	↑3.56	↑1.75	↑0.32
	animal	↑4.21	↑5.75	↑8.21	↑1.15	↑0.61
	color	↑3.82	↑4.61	↑7.06	↑1.74	↑0.54
Gemma	country	↑6.33	↑22.16	↓2.86	↑3.21	↓3.54
	animal	↑4.09	↑6.24	↑8.45	↑36.51	↓2.14
	color	↑4.65	↑5.16	↑7.81	↑16.49	↑1.21
OpenLlama	country	↑6.33	↑45.26	↑7.54	↑4.65	↓3.15
	animal	↑4.95	↑7.54	↑35.16	↑2.16	↑3.26
	color	↑4.39	↑5.27	↑27.56	↑11.42	↑2.51

Table 4: $\Delta \widehat{E}(t)$ compared to the previous token. The red represents $\widehat{E}(t)$ decreases compared to the previous token.

D ABLATION STUDY FOR SEMANTIC SENSITIVITY

To investigate the influence of different factors on IE value, we treat the performance of GPT2-XL on the “country” dataset as a baseline and implemented a series of variations. First, we replace GPT2-XL with other LMs, namely **model1** using GPT2-large, **model2** using GEMMA, and **model3** using OpenLlama. Second, we vary the dataset, forming **data1** using “animal” dataset and **data2** using “color” dataset. In addition, we use **candidate** to denote reduced candidates in the original “country” dataset (reducing the total number of countries from 25 to 15), and **fusion1**, **fusion2** to denote mixed candidates where fusion 1 mixes data from “country” and “animal” domains, and fusion 2 mixes data from “country”, “animal”, and “color” domains. Last, we alter the input sequence, forming **space** by replacing the *entity* in the 4th shot with *space+entity*⁸. **prefix** prepends a *comma* to the original first token.

As shown in Table 3, Differences in model IE become apparent only when a sufficient number of shots are provided. Statistically, models with larger parameter sizes exhibit higher IE. However, differences in data are apparent starting from the first shot, likely related to the domain of training data (token1, token2). Furthermore, depleting the diversity of shots effectively reduces the IE values in ICL (candidate). Lastly, the format of the prompt significantly influences IE, explaining LLMs’ sensitivity to certain perturbations (space, prefix).

E SUPPLEMENTARY MATERIALS FOR FINDING 1

E.1 LIMIT NUMBER OF SHOTS

In Section 3.3.1, we expanded the 3 categories into input sequences containing 10 shots (20 tokens each). Table 4 illustrates the changes in IE value for each shot relative to its predecessor within these sequences. The IE value of 4 different LLMs generally approached their upper limits by the 6th and 7th shots. It is important to note that these results only indicate the existence of an upper limit to the contribution of shot quantity to IE in ICL. They do not imply that the 6th and 7th shot universally represents the upper limit for all ICL tasks.

E.2 SHOT LENGTH

To examine the IE value associated with shot lengths, we designed a shot format pertinent to sentiment analysis as follows:

“*[entity]* sentiment: *[label]*,”

where “*[entity]*” represents emotional words such as “*happy*,” “*thrill*,” “*offended*”, etc., and “*[label]*” options include “*positive*” or “*negative*,” specifically chosen based on the category of

⁸In the GPT-2 tokenizer, *space+entity* is treated as a new token.

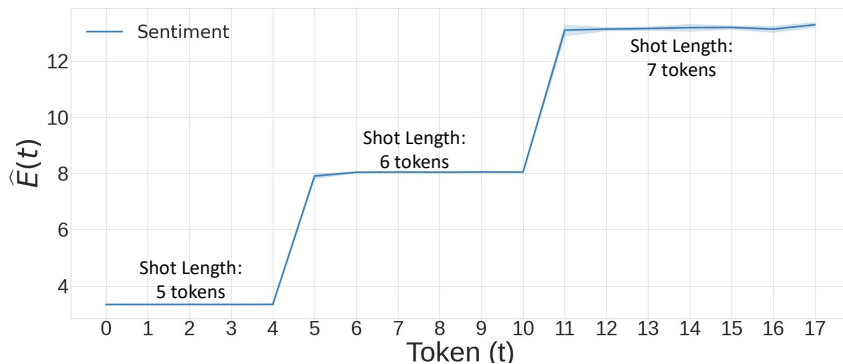


Figure 6: $\hat{E}(t)$ with inputs of 18 tokens, consisting of 3 shots in 5 tokens, 6 tokens, 7 tokens, respectively.

“[entity]”. The token length of “[entity]” was employed to control the overall length of the shot; for instance, when “[entity]” consists of single-token words like “anger,” “love,” etc., the entire shot spans 5 tokens, whereas for two-token words like “hopeful,” “resentful,” etc., the shot extends to 6 tokens. Consequently, we generated 300,000 input samples, each 18 tokens in length, comprising 3 shots with lengths of 5 tokens, 6 tokens, and 7 tokens respectively.

Figure 6 corroborates our hypothesis: within each shot, all tokens share a uniform IE value. This observation supports another intuitive viewpoint of ICL: an LLM gains greater confidence in the correctness of its predictions only when a new shot is introduced.

E.3 CASES OF INACCURATE GENERATIONS WITH EXCESSIVE SHOTS

In Table 1 we found 2 types of erroneous repetition, we listed some cases of them from GPT2-XL, in which blue text indicates the shots as prompt, the green text indicates correct entities and red text indicates wrong entities:

Case 1: The sequence breakdown was precipitated by the output of an incorrect entity.

“Ukraine, Mexico, Russia, Australia, New Zealand, United Kingdom, United States, Canada, United States of America, United States of America, United States of America, United States of America”

Case 2: Due to a loop spanning the shots, no new entities were generated.

“Canada, France, Turkey, Iran, Russia, Ukraine, United Kingdom, United States, Canada, Germany, United States, Canada, Germany, United States, Canada, Germany, United States, Canada, Germany”

F DETAILED MUTUAL INFORMATION TABLES

Tables 5 and 6 present the performance of GPT2-XL on the Animal category and OpenOrca datasets, respectively. Although “shots” and “natural sentences” demonstrate different patterns, they share a common characteristic: mutual information increases with token length, aligning well with the NTP mechanism.

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989

layer	token0	token1	token2	token3	token4	token5	token6	token7
1	2.83	2.83	6.89	6.50	10.68	9.24	14.16	11.53
2	2.83	2.83	6.90	6.91	11.08	11.10	16.70	16.79
3	2.83	2.83	6.89	6.88	11.17	11.17	16.93	16.00
4	2.83	2.83	6.89	6.88	11.08	11.06	16.74	16.58
5	2.83	2.83	6.89	6.89	11.13	11.11	16.94	15.88
6	2.83	2.83	6.88	6.89	11.16	11.16	18.89	17.11
7	2.84	2.83	6.89	6.88	11.15	11.14	16.97	17.08
8	2.83	2.83	6.88	6.88	11.12	11.19	16.86	17.42
9	2.83	2.83	6.90	6.88	11.20	11.17	16.92	15.97
10	2.83	2.83	6.88	6.88	11.08	11.14	16.97	16.51
11	2.83	2.83	6.88	6.88	11.04	11.05	17.05	16.19

990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

Table 5: Mutual information of GPT2-XL in Animal category. Red represents the highest value in this block.

layer	token0	token1	token2	token3	token4	token5	token6	token7
1	8.40	13.71	16.01	17.33	17.21	17.67	17.95	17.29
2	8.36	13.77	15.76	17.06	17.08	17.72	17.68	18.00
3	8.44	13.75	16.09	17.10	17.82	17.69	17.84	18.04
4	8.44	14.20	16.07	17.29	17.45	17.74	18.51	17.81
5	8.39	13.50	16.32	16.83	17.82	17.98	18.26	18.14
6	8.41	13.69	16.03	16.99	17.58	17.82	17.52	18.33
7	8.41	13.68	16.06	17.00	18.32	17.72	17.69	18.19
8	8.40	13.80	15.97	17.26	17.61	17.73	17.52	18.44
9	8.35	13.69	15.95	17.17	17.21	17.54	17.47	18.03
10	8.41	13.57	16.30	16.57	17.46	17.85	17.85	17.51
11	8.34	13.37	16.02	15.93	17.30	17.24	17.27	16.91

Table 6: Mutual information of GPT2-XL in OpenOrca dataset. Red represents the highest value in this block.