

Language Models Learn Universal Representations of Numbers and Here’s Why You Should Care

Michal Štefánik*^{1,2,3} Timothee Mickus*² Marek Kadlčík*³
Bertram Højer⁴ Michal Spiegel^{3,5} Raúl Vázquez² Aman Sinha⁶
Josef Kuchař³ Philipp Mondorf^{7,8} Pontus Stenetorp^{1,9}

¹R&D Centre for Large Language Models, National Institute of Informatics, Japan
²University of Helsinki ³TransformersClub @ Faculty of Informatics, Masaryk University
⁴IT University of Copenhagen ⁵Kempelen Institute of Information Technology
⁶IECL-ATILF, Université de Lorraine - ICANS Strasbourg
⁷MaiNLP, Center for Information and Language Processing, LMU Munich
⁸Munich Center for Machine Learning (MCML) ⁹University College London

*Equal contributors

Abstract

Prior work has shown that large language models (LLMs) often converge to accurate input embedding for numbers, based on sinusoidal representations. In this work, we quantify that these representations are in fact strikingly systematic, to the point of being almost perfectly universal: different LLM families develop equivalent sinusoidal structures, and number representations are broadly interchangeable in a large swathe of experimental setups. We show that properly factoring in this characteristic is crucial when it comes to assessing how accurately LLMs encode numeric and other ordinal information, and that mechanistically enhancing this sinusoidality can also lead to reductions of LLMs’ arithmetic errors.

1 Introduction

Mathematical problems provide an ideal playground in which to assess the current ability of large language models (LLMs) to manipulate precise information. While attempts to foster and improve mathematical reasoning in LLMs have been met with genuine empirical successes, e.g., as an aid to solving mathematical problems (Li et al., 2025), this field is still rife with challenges, particularly in ensuring that LLMs tackling mathematical problems remain robust and reliable. This calls for a deeper understanding of how robust the internal mechanisms of LLMs really are in operating over mathematical information.

Towards this goal, prior work identified that number representations stand out in terms of their sinusoidal structure (Nanda et al., 2023; Zhou et al., 2024; Levy and Geva, 2025; Kadlčík et al., 2025), accuracy of input embeddings (Kadlčík

et al., 2025), and systematicity of internal mechanisms in addition (Kantamneni and Tegmark, 2025; Zhou et al., 2024). While these works have provided pivotal insights as to how LLMs represent and manipulate numbers, their scope is limited: Kadlčík et al.’s findings only pertain to input embeddings; Zhou et al. find similarity of representations across models but do not quantify their alignment; Levy and Geva focus only on decimals as one particular aspect of representations. Crucially, *no* prior work considers *natural language* contexts, which strongly curtails the generalizability and practical usefulness of their findings, needed for the development of tools able to interpret – and verify – LLMs in practice.

To address these gaps and create a holistic picture of how language models handle numbers, this paper presents two main findings. First, we quantify that sinusoidal structures are **almost perfectly universal** across LLMs: models of different sizes and families converge to sinusoidal representations that are universal across *models* (§3.1) and *layers* (§3.3) while remaining highly systematic and *accurate* throughout the model’s internal processing (§3.2) – showing it is possible to track the processing of numeric information throughout models.

Second, we demonstrate the necessity of **suitable explainability tools** to make sense of how LLMs process numbers: we show how accounting for the sinusoidal nature of representations allows us to unravel the mechanism of manipulating large, multi-token numbers (§4.1) as well as other ordinal data (§4.3) in LLMs, and we mechanistically prove that a precision of sinusoidality in the model’s internal representation determines its output accuracy across arithmetic tasks (§4.2).

In summary, our main contributions are:

1. We develop a new, more accurate probe of numeric values that generalizes well across different tasks and models
2. We quantify that numeric representations of numbers are sinusoidal and highly accurate *throughout the model*, beyond the input embedding representations;
3. We are the first to analyze the LLMs’ numeric processing in *natural-language* settings and show that training numeric probes with real-world data is crucial for creating generalizing, practically useful probes;
4. We investigate mechanisms of numeric manipulations beyond addition and provide a mechanistic evidence that sinusoidality is crucial for models’ accuracy.

We make the new probe, with all our methods and analyses publicly available for any use.¹

2 Related work

Mathematical capabilities of LLMs The seminal findings of Brown et al. (2020) on remarkable arithmetical abilities of LLMs have ushered in a large body of work aimed at evaluating more complex mathematical capabilities of LLMs (Hendrycks et al., 2021; Cobbe et al., 2021; Sun et al., 2024; Yu et al., 2024) or lack thereof (Nikankin et al., 2025). This in turn has encouraged a focus on how to bolster the capabilities of LLMs or interpret how LLMs perform computations (Zhang et al., 2024; Stolfo et al., 2023). Most relevant to ours, some of the previous work focuses specifically on how numbers should be represented in principle. For instance, Charton (2022) assesses an impact on models performance in linear algebra problems when employing different encoding schemes based on scientific notation on linear algebra problems. Golkar et al. (2023) propose to encode numeric values by incorporating a scaled, learned control token <NUM>. Feng et al. (2024) remark that the precision of the numeric representation type impacts performance on arithmetic tasks, and argue to use the logarithmic-precision architecture of Feng et al. (2023).

Numeric representations in LLMs The increased focus on evaluating LLM performance

in verifiable domains such as math problems has also led to increased interest in the mechanisms by which LLMs compute arithmetic functions. Nanda et al. (2023) demonstrate how a model trained from scratch on modular addition relies on trigonometric operation: it maps inputs onto a unit circle, corresponding to a specific rotation, and then learns to combine the two rotations to derive a valid solution. Kantamneni and Tegmark (2025) apply circuit analysis to understand how general pretrained models perform addition, and report that representations are mapped onto a helix that can be manipulated using trigonometric operations. Zhou et al. (2024) refine these observations in qualitative analysis of Fourier components in LLMs’ representations (note that Fourier transformations map arbitrary functions into combinations of periodic trigonometric functions), and highlight the distinct role of attention and feedforward sublayers. Focusing more narrowly on representations rather than processing, Zhu et al. (2025) argue that non-linear probes do not provide a better fit than linear probes. Levy and Geva (2025) remark that it is possible to retrieve the digit (value mod 10) of numeric inputs with high accuracy. Our work builds upon Kadlčik et al. (2025) that propose a probe architecture that factors in the sinusoidal nature of representations, and show that they can retrieve numeric values from input embeddings with much higher accuracy.

3 Universality of Sinusoidal Representations of Numbers

3.1 Different Models Learn Equivalent Representations of Numbers

The first point we address is the distinctiveness of LLMs’ representations of numbers. Results from Kadlčik et al. (2025) and Zhu et al. (2025) suggest that different models converge to the same *type* of number representations, but they do not provide a direct assessment of how *closely* the representations match across models – which could evidence that a shared representation is a causal consequence of architectural bias and optimization process rather than a coincidental artifact. Following Kadlčik et al., we focus on the input embeddings from eight LLMs of diverse sizes and families with open-sourced, pre-trained checkpoints: OLMo 2 (OLMo et al., 2025), Llama 3 (Grattafiori et al., 2024) and Phi 4 (Abdin et al., 2024).

To quantify whether models converge to similar embeddings, we start by conducting a simple Rep-

¹<https://github.com/promptheus/numllama>

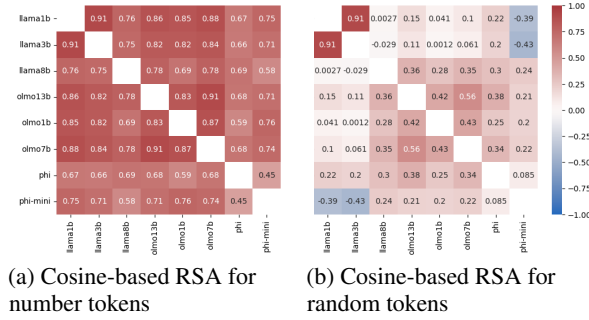


Figure 1: RSA scores for input embeddings of LLMs

representational Similarity Analysis (RSA; [Kriegeskorte et al., 2008](#)) across the input embeddings for number word-pieces in paired models. As a baseline, we compute RSA scores for a random sample of 1000 pieces present in the vocabularies of all the models under consideration. Results of the analysis are displayed in Figure 1. Across all pairs of models, we find that number embeddings systematically yield consistently higher RSA scores than a random sample of word pieces. While some models (especially in the Phi family) are not as well aligned with other models, we do observe high scores both within and across families, demonstrating that models converge to embedding spaces with **equivalent similarity structure**.

Another approach to quantifying the similarity of number representations is to inspect their base frequencies through Fourier decompositions ([Kadlčík et al., 2025; Zhou et al., 2024](#)). Applying a PCA transformation followed by a Fourier transform allows us to quantify the *magnitude* of individual frequencies (Figure 2a). Subsequently, we determine whether two models agree as to how they rank frequencies by computing a simple intersection-over-union (IoU) of the top k frequencies. We repeat this process for every pair of models, using the input embeddings of number pieces (Figure 2b) as well as that of random pieces (Figure 2c). When considering the top $k = 63$ frequencies, we find **perfect agreement across all models** for numbers. This provides a quantitative basis for some of [Zhou et al.’s \(2024; §C.1\)](#) qualitative observations.²

In summary, our analyses show that the LLMs from different families learn to represent numbers in the same topology and trigonometric features using the same dominant frequencies. This is specific to numbers: we do not reproduce these observa-

²See our Section A.1 for technical details, including how the value of k is set.

tions with other word pieces shared across models.

3.2 Number Representations are Always Sinusoidal

In Section 3.1, we presented evidence that the input embeddings of numbers *across* different LLMs are highly consistent. We now focus on whether the sinusoidal mechanism of representations is maintained and employed *throughout* the model.

For qualitative assessment, we visualize the PCA and its Fourier transform of internal activations of numerical tokens in Section A.4. Besides the apparent wave-like pattern, the representations have a sparse Fourier transform, confirming the sinusoidal character. For quantitative analysis, we look at the model’s internal representations through the lens of the sinusoidal probe proposed by [Kadlčík et al. \(2025\)](#). This probe was designed to map input embeddings of LLMs into an integer, thus classifying embeddings into predefined range of numeric values. It is defined as:

$$f_{\sin}(\mathbf{x}) = (\mathbf{W}_{\text{out}}\mathbf{S})^T(\mathbf{W}_{\text{in}}\mathbf{x})$$

$$\mathbf{S}_{ij} = \begin{cases} \sin(ie^j 1000/h) & \text{if } j \equiv 0 \pmod{2} \\ \cos(ie^{j+1} 1000/h) & \text{if } j \equiv 1 \pmod{2} \end{cases} \quad (1)$$

where $\mathbf{W}_{\text{in}} : h \times d$ and $\mathbf{W}_{\text{out}} : h \times d$ are learned parameters, and $\mathbf{S} : h \times 1000$ injects an inductive bias in the classifier towards sinusoidality.³


While [Kadlčík et al.](#) demonstrate that such a probe properly benefits from the expected sinusoidal inductive bias and greatly bolster accuracy for number representations, their approach does not enforce a strict sinusoidal structure and may fail to generalize due to redundancy in a model’s representation. As an alternative, we design a new probe which assigns a *single sine wave to each latent feature* and *directly* optimize the frequency, phase-shift, and amplitude of each sine wave:

$$f_{\text{param-sin}}(\mathbf{x}) = \mathbf{S}^T(\mathbf{W}_{\text{in}}\mathbf{x})$$

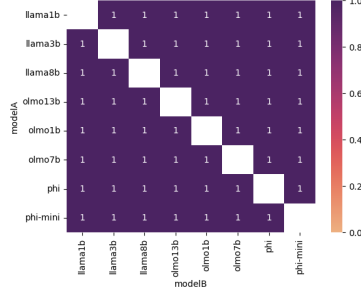
$$\mathbf{S} = \mathbf{w}_a \sin(\mathbf{w}_\phi + \mathbf{w}_f \mathbf{r}^T) \quad (2)$$

where \mathbf{r} are constant label values (0...999 for all models we analyze), and $\mathbf{w}_f, \mathbf{w}_a, \mathbf{w}_\phi \in \mathbb{R}^d$ are three vectors of learnable parameters modeling the frequencies, amplitudes and the phase-shifts of the sine waves. This architecture *strictly constrains* the classification matrix \mathbf{S} to take the form of d sine waves sampled at n evenly-spaced points. The

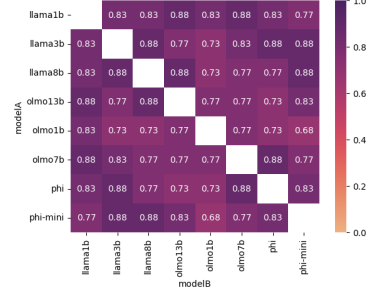
³Unless otherwise stated, we use $d = 100$; h corresponds to the inner dimensionality of the LLM at hand.



(a) Contributions of Fourier base frequencies (shown example: Llama-3 1B)



(b) IoU for number pieces



(c) IoU for random pieces


Figure 2: Intersection-over-union of top $k = 63$ Fourier base frequencies

configuration of initialization and hyperparameters is provided in Appendix B.


We first assess whether sinusoidal probes are the most suitable, i.e., an accurate choice for decoding the *internal* representations of the model. To contextualize the accuracy in terms of sinusoidal quality of the representation, we also evaluate other types of probes used to decode numbers in previous work (Feng et al., 2023; Zhu et al., 2025; Kadlčik et al., 2025) as baselines. We fit separate probes for activations from every layer across six different language models from three families. Contrary to previous work, we assess the representations of numbers in *natural-language* contexts, rather than as single tokens or synthetic arithmetical commands. Our contexts cover four domains (temporal, arithmetic, medical and culinary) covering numerical contexts including quantities, measurements or timestamps (detailed in Appendix A.2). To assess generalization rather than memorization capacity, we *split* the prompts and corresponding representations into training, validation and test sets containing *distinct* extracted numbers (x_2).

Figure 3a shows the accuracy of different probes extracting the value of input numbers from representations of each model’s layer, evaluated across six different models. As we can see, both sinusoidal probes reach over 80% of accuracy in all but two cases and over 90% of accuracy for a majority of probing scenarios. Here, the superior accuracy of both sinusoidal probes provides an indicator of both (i) the sinusoidal character, and (ii) the accuracy of input number representation across layers — the best probe results present a lower bar of the accuracy of numeric representation on a given layer.

Figure 3b disentangles the accuracies of the param-sin probe from Figure 3a by the model, covering six LLMs of different sizes and three different



(a) using probes from prior work with different bases



(b) across different language models; param-sin probe

Figure 3: **Accuracy of decoding numbers** from internal activations of language models

families. While the probing is slightly less accurate on some particular layers of particular models, for none of the analyzed models does the sinusoidal probe systematically underperform. If we take into account that our probes need not be perfect and reflect a *lower bound* of real accuracy of numeric representations, these results taken together show that a sinusoidal representation is a decisively dominant scheme of numeric representations in language models, corroborating and extending the observations of Kantamneni and Tegmark (2025).

A vast majority of cases performing lower than 80% occurs in probing the models’ first and a last layer; in Section A.6, we analyze these cases in detail through the lens of probes learned representations and find the substantial qualitative differ-

ence in the *sparsity* of representations of numbers. While the middle layers of all models resort to a ‘dense’ representation decomposed into a small set of features bearing full numeric information, the initial and final layers scatter the numeric information across most of the latent features. As the scattered features share the information with natural language, a robust recovery in a realistic setting of natural contexts becomes more difficult.

Generalization across contexts We can see that sinusoidal probes trained with natural-language contexts can provide highly accurate estimates of LLMs’ numerical states across models. However, favoring simplicity, all of the previous work (§2) investigated numerical representations in *synthetic* mathematical contexts. To assess whether such a simplification comes for a price in practical applicability of resulting probes, we explore whether the choice of training contexts affects generalization of resulting probes.




Figure 4: **Generalization of probes** fitted on representations in natural language (yellow), and synthetic, mathematical contexts (blue), evaluated on natural-language (solid) and math contexts (dashed). Llama 3.2 1B.

In Figure 4, we compare the performance of probes trained on natural-language contexts (yellow) and synthetically-constructed math contexts (blue), evaluated on natural-language inputs (solid lines) and synthetic inputs (dashed lines). We can see that probes fitted with natural-language contexts are much more robust in application across *both* math and natural-language context. While both probes generalize well within their training context type, **only the natural-language probes generalize across evaluation settings**, robust to the context’s distribution shift. This observation has two important implications. First, it informs future work aiming to explain models’ internal computations in practical applications to favor probes trained in natural-language settings. Second, it restrains future work in interpretability and explain-

ability from drawing general conclusions on models’ mechanics from fully-synthetic settings, evidencing that such conclusions may not generalize to real-world applications.

In summary, our results underscore that the same type of sinusoidal representation of numbers holds in general across different model types. Numbers are represented in a similar, systematic fashion, regardless of which layer or type of context we consider. While we can expect this to be in part due to the residual stream across layers, we note that there is no a priori reason for an LLM to preserve sinusoidality across layers, and previous work shows that input tokens tend to have a negligible impact on upper layer representations (Mickus et al., 2022).

3.3 Different Layers use Interchangeable Numeric Representations

Having found the remarkable degree to which the embeddings for numbers align across different models and the noteworthy universality of sinusoidal representations in varied contexts, as a final step, we turn to assess whether the models utilize the identical representation throughout its computation. To investigate that, we employ the probes from Section 3.2 trained separately for each separate model layer (L_i) and measure the accuracy of each probe on representations derived for every *other* layer $L_j \neq L_i$.

While informative, this approach still comes with the caveat that probes trained for a given layer L_i may pick up on idiosyncrasies inherent to a specific layer — i.e., the probes might not disentangle what is specific to numbers as opposed to what is specific to L_i . To address this point, we also fit probes using *all but one* layer ($L_1, L_2, \dots, L_{i-1}, L_{i+1}, \dots, L_n$), and evaluate performance on representations from the held out layer L_i . As before, we hold out 100 numbers for assessing generalization in validation and test conditions.

Results of probes’ cross-layer generalization for Llama 1B are displayed in Figure 5. We observe that probes fitted on a concrete layer’s representations (Figure 5a) generalize outstandingly well to both close-by and more distant layers, with the exception of the first and the last layer. This trend reveals that language models **operate in a consistent representation** of numbers that is **universal throughout their computation**, undergoing only minor shifts across layers. This trend is also corroborated by multi-layer probes evaluated in heldout-layer fashion (Figure 5b), reaching a close-to 100%

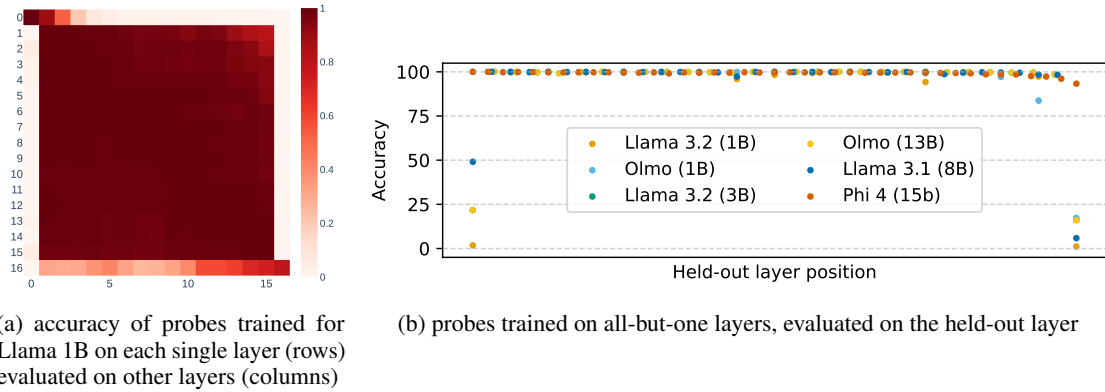


Figure 5: **Transferability of representations:** Accuracy of param-sin probe evaluated on **unseen layers**

accuracy across *all* intermediate layers and six models of diverse sizes and families. This result has practical implications for future work in numeric interpretability — showing that we can train universal, yet highly *accurate* probes for intermediate layers of a broad set of language models.

To summarize, we find that the same key, i.e. probe can be used to interpret numeric information at different layers of the same language model. Added to our previous observations in Sections 3.1 and 3.2, we can therefore stress that models learn input embeddings for numbers that are **strikingly similar**, regardless of which model they come from. These embeddings are processed into sinusoidal representations that are systematic and consistent regardless of context or hidden layer, with the source of this consistency maintained primarily by the residual stream across layers (cf. Section A.7).

4 How to Study Numbers in LLMs

We have established that models universally converge to sinusoidal number representations. This characteristic is not shared with other word pieces. This questions the appropriateness of existing interpretability tools when it comes to numbers and other ordinal data types.




Figure 6: Accuracy of extracting pieces of **multi-token numbers** from the last token representations.

4.1 Multi-token Numbers

As a case study for observing how the choice of interpretability toolset affects our findings, we analyse whether our conclusions about the remarkable accuracy of numeric representations also hold for large, *multi-token* numbers.

First, we use our sinusoidal probes to retrieve the values of *all parts* of multi-token numbers from a *single*, i.e. last token’s representation. For this experiment, we adapt the methodology utilizing natural-language contexts (see Figure 4), with a substitution of original numeric values with values spanning two to six numeric tokens, i.e. in a numeric range between $10^3 - 10^{18}$. We use a subset of BBC data that was published after the cutoff or release date of our LLMs (Li et al., 2024).

As a point of comparison, we consider the general-purpose TunedLens tool (Belrose et al., 2025), an early-exit interpretability method which aims to explain the contents of intermediate representations by learning to translate them into logits. For each layer L_i , TunedLens distills the computations done in the subsequent layers L_{i+1}, \dots, L_n into a simple affine transformation \mathbf{W}_i , such that applying this transformation \mathbf{W}_i to a hidden state at L_i followed by the unembedding matrix closely matches the logits that the LLM would eventually produce. TunedLens was applied to number representations in previous work (Zhou et al., 2024) and also represents the line of work arguing for a linear representations of numeric values in LLMs (Zhu et al., 2025). To maximize comparability and align the number and distribution of target categories, we narrow down the space of probed tokens to 1,000 tokens ranked as 2,000–3,000th most-common tokens in our dataset (matching the number of classes and their distribution for our number probes). Crucially, we also report previous-token recovery rates using

TunedLens for *number* tokens, using the same configuration as for natural-language tokens.

Results for Llama 1B and different probing methods are presented in Figure 6. Results of sin probes show that **multi-token numbers are indeed systematically and highly accurately superposed in the latest numeric representation** – reaching an accuracy of 99% for the immediately-preceding number piece (offset -1). However, the accuracy and/or systematicity of the superposition mechanism quickly drops for numbers longer than *three* tokens (i.e. $\geq 10^9$), reaching close to zero for the 5th preceding token. More central to our point, **TunedLens gives a conflicting account**. TunedLens suggests that non-number tokens can be decoded as far as five time-steps later with non-zero probability, but that number tokens cannot be retrieved past two time-steps, and even then, only with 50% accuracy. This comparison shows that TunedLens is highly inefficient for recovering numeric tokens compared to sinusoidal probes; estimates derived from it provide an inaccurate representation of LLMs’ capabilities, underestimating the amount of contained information. A broader point to be made is that this suggests a trade-off between the coverage of an interpretability tool and its appropriateness to particular situations.

4.2 Output Tracking

A reliable mechanism of numeric manipulations necessitates an accurate representation of inputs (evidenced in §3), but also the ability to construct systematic and accurate representations of *outputs*. Without the specialized tools to analyze numeric representations, prior work describes LLMs’ arithmetical mechanisms through qualitative assessments limited to the outputs of addition (Kantamneni and Tegmark, 2025; Zhou et al., 2024; Levy and Geva, 2025). Therefore, having found the sinusoidal probes to be highly suitable for describing the representations of input numbers, we set off to investigate whether such probes can also better capture the mechanics of LLMs’ *output* representations, compared to other probe types.

In this set of experiments, we prompt the language models with prompts of the form “*The result of $x_1 \otimes x_2$ is*”⁴ and train probes to predict the correct result of $x_1 \otimes x_2$, where $\otimes \in \{+, -, \times, /\}$.

Figure 7 shows the results for probes from

⁴Exact prompt formulations for each operation can be found in the supplementary material under /notebooks/output_probing.

prior work, including linear probes from Zhu et al. (2025); Zhou et al. (2024) and a sin probe proposed and applied by Kadlčík et al. (2025). We can see that the param-sin probe performs comparable or better than probes applied for numbers in previous work and newly allows the tracking of outputs with over 80% of accuracy in multiplication. The results for addition also push the known lower bounds of precision of the result emerging early in the model, conflicting with a hypothesis of “gradually refining predictions” from Zhou et al. (2024) based on LogitLens, but supporting the interpretation of arithmetics implementation in LLMs by Stolfo et al. (2023) through a combination in late layer(s).

Figure 8 further details on a distinction between the cases where the model predicts a correct and incorrect response – suggesting that the **extent to which the model maintains the sinusoidal representation may determine its accuracy**.

We evidence this by following qualitative and mechanistic evidence. In Section A.5, we observe the output embeddings of addition and multiplication – latter coupled with lower model’s accuracy – differ in the precision of the the observed sinusoidality; with the pattern being more prominent in the better-performing, addition embeddings. In Section A.8, we further analyze how, according to our probes, is the result refined via gradually minimizing the response error across layers, with particular layers deteriorating this refinement, suggesting a worse fit of the sinusoidal probe and thus, the sinusoidal pattern. We find that simply dropping the layers causing a divergence from the probe’s fitted pattern lead to an accuracy improvement in multiplication and division in 4 out of 6 cases.

Finally, we evidence this point more directly by a mechanistic intervention – **steering activations towards the expected sinusoidal pattern**. Concretely, we (i) fit sinusoidal probes at each layer to predict the output of a given arithmetic operation; (ii) optimize a set of 1,000 randomly-initialized ‘patch’ embeddings e_y to maximize the probability that the probes will map them onto the integer value y ; and (iii) steer activations towards the corresponding patch whenever the model is not producing the expected output. Steering is achieved by interpolating between the activation h_i at layer L_i and the patch e_y for the target scalar y , $\alpha h_i + (1 - \alpha)e_y$. We report results with respect to the optimal value of α per layer, operation and probe.

Figure 9 displays the absolute difference in per-

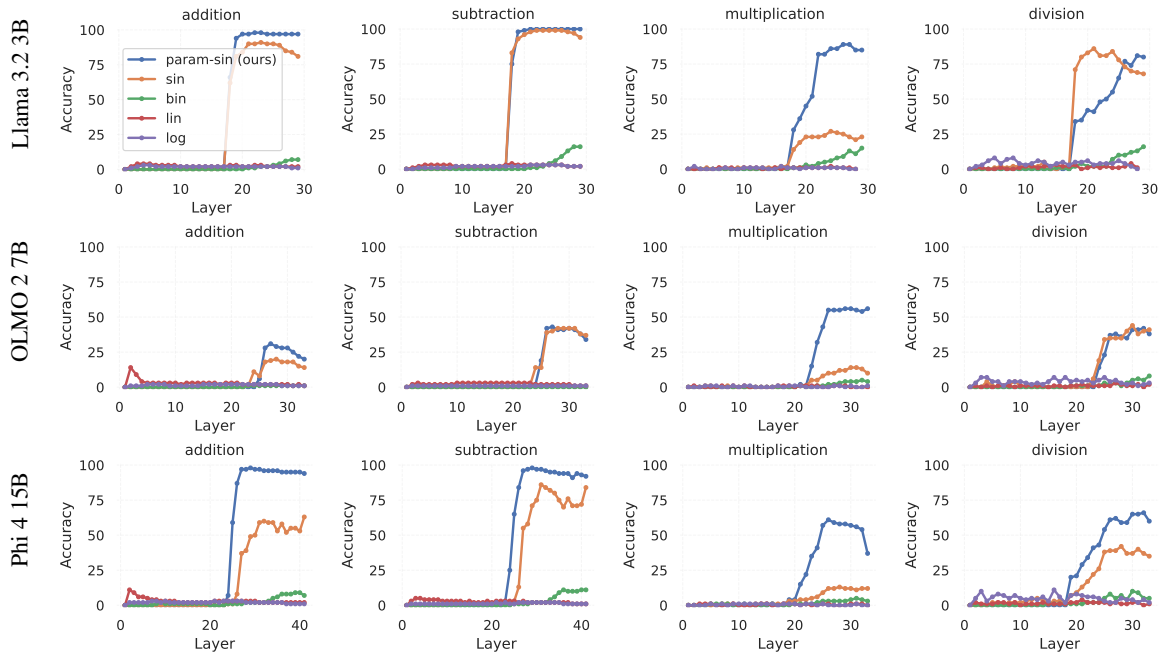


Figure 7: **Probing accuracy** in extracting correct output from internal activations across layers; different probes.

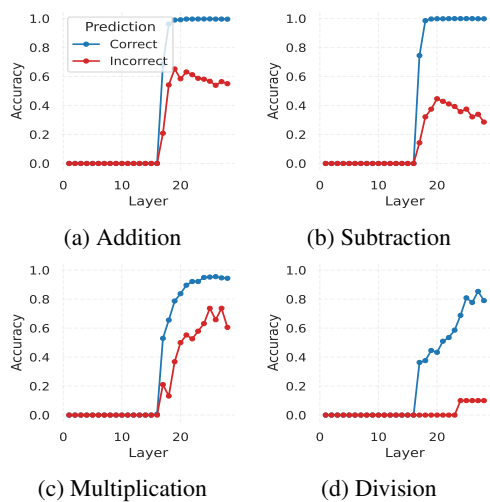


Figure 8: **Probing accuracy** in extracting correct output from Llama 3.2 3B’s internal activations separates model’s correct (blue) and incorrect (red) predictions.

formance caused by the steering on previously erroneous cases in multiplication and division for three models. We target probes on upper layers which reach accuracy scores $\geq 10\%$. In short, depending on the layer, we can correct up to 30% of errors for multiplication, and up to 42% for division. Crucially, our probe almost always outperform the probe proposed by Kadlčík et al. (2025): Across all setups, we only observe four setups where the sin probe outperforms our sin-param probe.⁵ The

⁵Viz., division with OLMo at L22, L23, L24), multiplication with Phi at L37.

optimal setup per model and operation is always obtained with our probe. This underscores once again the importance of selecting the right explanatory tool, this time in the context of activation steering.

In sum, factoring in the knowledge of the sinusoidal nature of number representations allowed us to understand and assess the lower bounds of accuracy of multi-token number representations. On a case study of Llama 3.2 3B, we further showed how more accurate probes allow pinpointing and **eliminating** sources of errors. Tools that are best suited to number representations allow us to enhance accuracy and perform more effective interventions.

4.3 Ordinal Data Types

Having established that sinusoidal mechanisms robustly encode numeric values and our specialized probe facilitates targeted interventions, we investigate how this tooling generalizes to ordinal data types – potentially represented in the same systematic fashion as numbers. A particular challenge here is a data scarcity – for instance, as there are only seven days in a week, to warrant held-out evaluation, we can only fit the probes on six examples.

Prior work by Engels et al. (2025) demonstrates that LLMs internally represent temporal data – such as days of the week and months of the year in circular structures. Analytically, a centered circle is parameterized by the equations $x(\theta) = r \sin(\theta + \pi/2)$ and $y(\theta) = r \sin(\theta)$, indicating that these temporal concepts are encoded using sinusoidal features.

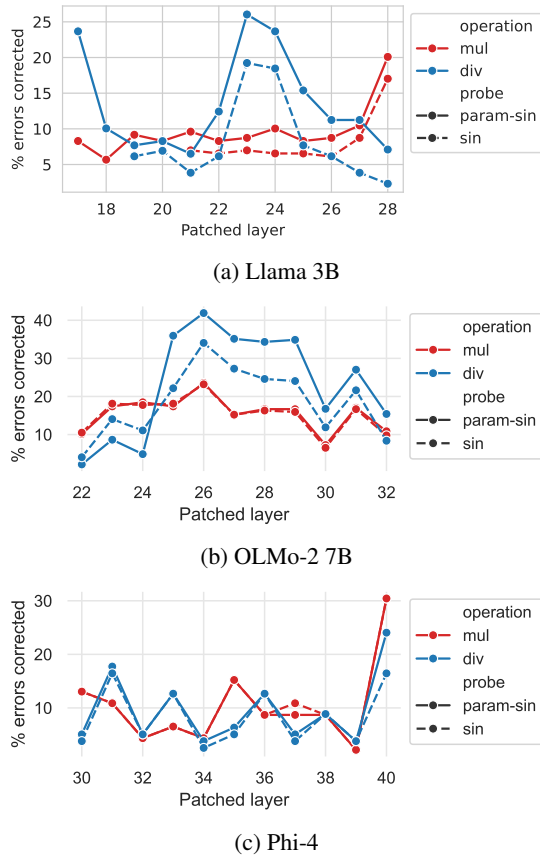


Figure 9: **Error reduction** after **steering activations** on incorrect predictions of multiplication and division. Optimal interpolation factor α per setup.

Assuming this structure, we fit our probe on the embedding representations of these concepts and compare it to a prior probe architecture of Kadlčík et al. (2025) in leave-one-out cross-validation setting. The best held-out validation accuracies of both architectures are reported in Table 1. Our results demonstrate high decoding accuracy for ordinal data across all evaluated models. This validates the probe’s utility beyond pure numbers, even under extreme data-scarcity. With our better generalizing architecture, we corroborate and extend the discovery by Engels et al. (2025) of circularity in LLMs’ representations to other ordinal data types.

5 Conclusions

In this paper, we have shown the number representations in LLMs are exceptional in two distinct but related ways. First, the regularity and systematicity of the sinusoidal structure that all LLMs converge is striking, to the point of being almost universal across all experimental setups we consider. We can view the same frequencies in Fourier being used by all models we survey (§3.1), we can train a probe

Model	sin	param-sin
<i>Days of week (n=7)</i>		
Olmo 2 1B	4	6
Olmo 2 7B	6	6
Olmo 2 13B	5	6
Llama 3 1B	5	6
Llama 3 3B	7	6
Llama 3 8B	3	6
<i>Months (n=12)</i>		
Olmo 2 1B	6	11
Olmo 2 7B	6	11
Olmo 2 13B	5	11
Llama 3 1B	9	11
Llama 3 3B	10	11
Llama 3 8B	6	11
<i>Numerals: ‘one’ to ‘ten’ (n=10)</i>		
Olmo 2 1B	3	6
Olmo 2 7B	4	8
Olmo 2 13B	4	8
Llama 3 1B	3	5
Llama 3 3B	4	6
Llama 3 8B	3	9

Table 1: Number of matches achieved by training param-sin and best-performing prior probe (sin) in decoding ordinal data with extremely small training sets from models’ internals. Held-one-out evaluations.

on one layer and apply to another (§3.3). The one remarkable exception is the discrepancy between natural language contexts, which yield remarkably robust probes that generalize broadly, and synthetic probes derived from templated texts, which do not.

Second, this exceptional sinusoidal regularity requires adapting the way we study number processing in LLMs. Using a general-purpose TunedLens method, applied for numbers in prior work, grossly underestimates how accurately previous timestep information can be retrieved, whereas sinusoidal probes such as the one we propose in eq. (2) may lead to substantially higher accuracies and contradictory conclusions (§4.1). Finally, we demonstrate that the sinusoidal probes are also well applicable for analyzing the emergence and manipulation of model’s own predictions and that the accuracy of representation in terms of sinusoidality is related to the prediction accuracy (§4.2).

While we focus on numbers, our experiments showcase how important the appropriate interpretability toolset is for accurate interpretation of specialized data types. This naturally questions the common practice of designing *general-purpose* interpretability and explainability tools, motivating future work towards adapting our practices to the specifics of the type of data we want to explain.

Limitations


Several practical factors limit the scope of our work. We focus our survey on three particular model families; moreover, we have excluded very large models (>30B parameters) due to practical hardware limitations. The rapid pace of modern NLP may entail that the set of models we have picked for this study will no longer be representative of contemporary frontier models in the future.

Our work is also based on English-centric LLMs, but the number systems vary greatly across languages — e.g., Chinese and Japanese often employ myriad-based number systems and frequently use both Chinese and Arabic numerals. The extent to which our conclusions will hold for a Chinese-centric or Japanese-centric LLM remains a topic for future research.

Lastly, we find that all errors in probing temporal concepts (days, months) are exactly in the first position (Monday, January), persisting even when the probed token list is reversed (Sunday, December). We suspect that zero value is a blind spot of the probe due to its parametrization, which needs to be addressed to achieve a perfect decoding rate.

Acknowledgments

This work is supported by the Research Council of Finland through project No. 353164 “Green NLP – controlling the carbon footprint in sustainable language technology”.

 This project has received funding from the European Union’s (EU) Horizon Europe research and innovation programme under Grant agreement No. 101070350. The work has also received funding from the Digital Europe Programme under grant agreement No 101195233 (OpenEuroLLM). The contents of this publication are the sole responsibility of its authors and do not necessarily reflect the opinion of the EU.

References

Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, James R. Lee, Yin Tat Lee, Yuanzhi Li, Weishung Liu, Caio C. T. Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli Saarikivi, and 8 others. 2024. [Phi-4 technical report](#). *Preprint*, arXiv:2412.08905.

Nora Belrose, Igor Ostrovsky, Lev McKinney, Zach Furman, Logan Smith, Danny Halawi, Stella Biderman,

and Jacob Steinhardt. 2025. [Eliciting latent predictions from transformers with the tuned lens](#). *Preprint*, arXiv:2303.08112.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020. [Language models are few-shot learners](#). *Preprint*, arXiv:2005.14165.

Francois Charton. 2022. [Linear algebra with transformers](#). *Transactions on Machine Learning Research*.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.

Joshua Engels, Eric J Michaud, Isaac Liao, Wes Gurnee, and Max Tegmark. 2025. [Not all language model features are one-dimensionally linear](#). In *The Thirteenth International Conference on Learning Representations*.

Guhaio Feng, Kai Yang, Yuntian Gu, Xinyue Ai, Shengjie Luo, Jiacheng Sun, Di He, Zhenguo Li, and Liwei Wang. 2024. [How numerical precision affects mathematical reasoning capabilities of llms](#). *Preprint*, arXiv:2410.13857.

Guhaio Feng, Bohang Zhang, Yuntian Gu, Haotian Ye, Di He, and Liwei Wang. 2023. [Towards revealing the mystery behind chain of thought: A theoretical perspective](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.

Siavash Golkar, Mariel Pettee, Michael Eickenberg, Alberto Bietti, Miles Cranmer, Geraud Krawezik, Francois Lanasse, Michael McCabe, Ruben Ohana, Liam Parker, Bruno Régaldó-Saint Blancard, Tiberiu Tesileanu, Kyunghyun Cho, and Shirley Ho. 2023. [xval: A continuous number encoding for large language models](#). In *NeurIPS 2023 AI for Science Workshop*.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. [Measuring mathematical problem solving with the MATH dataset](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.

- Marek Kadlčík, Michal Štefánik, Timothee Mickus, Josef Kuchař, and Michal Spiegel. 2025. [Pre-trained language models learn remarkably accurate representations of numbers](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 26693–26702, Suzhou, China. Association for Computational Linguistics.
- Subhash Kantamneni and Max Tegmark. 2025. [Language models use trigonometry to do addition](#). *Preprint*, arXiv:2502.00873.
- Nikolaus Kriegeskorte, Marieke Mur, and Peter Bannettini. 2008. Representational similarity analysis - connecting the branches of systems neuroscience. *Front. Syst. Neurosci.*, 2:4.
- Amit Arnold Levy and Mor Geva. 2025. [Language models encode numbers using digit representations in base 10](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 385–395, Albuquerque, New Mexico. Association for Computational Linguistics.
- Yucheng Li, Frank Guerin, and Chenghua Lin. 2024. [Latesteval: Addressing data contamination in language model evaluation through dynamic and time-sensitive test construction](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(17):18600–18607.
- Zenan Li, Zhaoyu Li, Wen Tang, Xian Zhang, Yuan Yao, Xujie Si, Fan Yang, Kaiyu Yang, and Xiaoxing Ma. 2025. [Proving olympiad inequalities by synergizing llms and symbolic reasoning](#). *Preprint*, arXiv:2502.13834.
- Timothee Mickus, Denis Paperno, and Mathieu Constant. 2022. [How to dissect a Muppet: The structure of transformer embedding spaces](#). *Transactions of the Association for Computational Linguistics*, 10:981–996.
- Neel Nanda, Lawrence Chan, Tom Lieberum, Jess Smith, and Jacob Steinhardt. 2023. [Progress measures for grokking via mechanistic interpretability](#). In *The Eleventh International Conference on Learning Representations*.
- Yaniv Nikankin, Anja Reusch, Aaron Mueller, and Yonatan Belinkov. 2025. [Arithmetic without algorithms: Language models solve math with a bag of heuristics](#). *Preprint*, arXiv:2410.21272.
- Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, Nathan Lambert, Dustin Schwenk, Oyvind Tafjord, Taira Anderson, David Atkinson, Faeze Brahman, Christopher Clark, Pradeep Dasigi, Nouha Dziri, and 21 others. 2025. [olmo 2 furious](#). *Preprint*, arXiv:2501.00656.
- Alessandro Stolfo, Yonatan Belinkov, and Mrinmaya Sachan. 2023. [A mechanistic interpretation of arithmetic reasoning in language models using causal mediation analysis](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7035–7052, Singapore. Association for Computational Linguistics.
- YuHong Sun, Zhangyue Yin, Qipeng Guo, Jiawen Wu, Xipeng Qiu, and Hui Zhao. 2024. [Benchmarking hallucination in large language models based on unanswerable math word problem](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2178–2188, Torino, Italia. ELRA and ICCL.
- Longhui Yu, Weisen Jiang, Han Shi, Jincheng YU, Zhengying Liu, Yu Zhang, James Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. 2024. [Metamath: Bootstrap your own mathematical questions for large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Wei Zhang, Chaoqun Wan, Yonggang Zhang, Yiu ming Cheung, Xinmei Tian, Xu Shen, and Jieping Ye. 2024. [Interpreting and improving large language models in arithmetic calculation](#). *Preprint*, arXiv:2409.01659.
- Tianyi Zhou, Deqing Fu, Vatsal Sharan, and Robin Jia. 2024. [Pre-trained large language models use fourier features to compute addition](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 25120–25151. Curran Associates, Inc.
- Fangwei Zhu, Damai Dai, and Zhifang Sui. 2025. [Language models encode the value of numbers linearly](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 693–709, Abu Dhabi, UAE. Association for Computational Linguistics.

A Supplementary results

A.1 Selecting the optimal number of Fourier components for comparison

In Section 3.1, we discuss selecting the optimal number of Fourier components. As can be assessed from 2a, from a certain frequency rank, Fourier decompositions resort to noise. Thus, our objective is to find the “cutoff” rank (k) that disentangles oscillating information from the non-oscillating, presumably non-numeric information. This means finding an *optimal*, yet non-trivial (e.g. ≥ 10) cutoff for a number of components. In practice, we evaluate for each value of k what the minimum IoU agreement score across all models amounts to for numbers and random overlapping word-pieces. As displayed in Figure 10, we find such an optimal cutoff at $k = 63$ as the highest value of k that leads to perfect agreement across models.

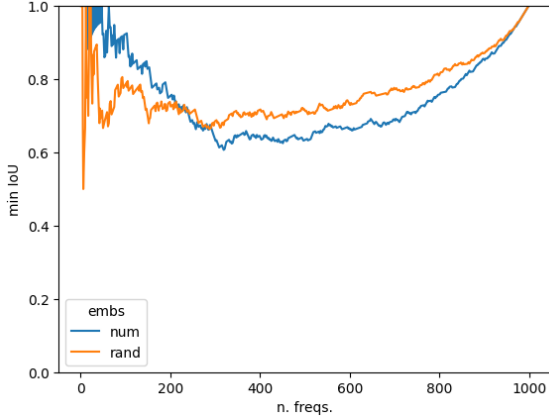


Figure 10: Minimum IoU across number of frequencies considered

A more thorough assessment of Figure 10 suggests a few interesting trends. Random word pieces also tend to favor a handful of Fourier basis components; which we conjecture is due to the sampling mechanism. By selecting overlapping pieces, our pieces must be frequent enough to be present in multiple distinct tokenizers, which in turns shapes the type of linguistic units represented in this random sample. Secondly, we also observe a ‘cross-over’ point around $k \approx 250$, after which we find greater agreement in random pieces than numbers. Yet, we note that (i) a significant proportion of the mass is concentrated in a few frequency components for numbers, whereas random pieces lead to much more uniform distributions across frequency components (see Figure 11); and (ii) the results in Figure 10 still allow us to establish that the Fourier profile of number pieces is clearly distinct from what we observe for any other overlapping set of word pieces.




Figure 11: Contributions of Fourier base frequencies for non-number pieces (Llama-3 1B, provided as an example)

	Add	Sub	Mul	Div
Accuracy	100%	99.8%	90.4%	5.9%
P(Extracted Incorrect)	-	56.8%	26.3%	94.4%
P(Not extracted Correct)	1.4%	0.03%	26.1%	5.9%

Table 2: **Probe accuracy in extracting the predicted result:** Ratio of cases where the sin probe of some layer (top) retrieves a correct result when the model’s prediction is *incorrect*, and (bottom) can *not* retrieve a correct result when the model’s prediction is *correct*.

A.2 Natural-language contexts for representations extraction

Contrary to prior work, we use the contexts of natural language to extract the representations of numbers. We argue that such setting better corresponds to the practical applicability of probes – in monitoring models’ behavior and errors, or in tracking the faithfulness of reasoning chains.

We choose the contexts with medium to high co-dependency on numerical contents covering four distinct domains and eight different datasets. The datasets categorized into these domains are presented in Table 3 and total in around 700 thousands natural-language contexts.

A.3 Generalization of probes to unseen layers

Figure 12 shows how probes fitted on each layer generalize to all other layers in a model, using synthetic mathematical contexts. Strong cross-layer generalization indicates a high consistency of representations throughout the model.

A.4 Visualizations of internal activations

We visualize the internal activations for a string template “ $x1 + x2 =$ ” on the second numeric token ($x2$). Values for $x2$ are selected as a range 0-999 (all values present by the vocab), and values for $x1$ are sampled randomly from the same range. Then, we project the activations of the model’s middle layer to 64 dimensions with PCA and compute the Fourier transform.

We visualize the first 16 PCA dimensions in Figure 13 and the maximal magnitudes of the frequencies in the Fourier transform in Figure 14. We visualize across model sizes and families.

A.5 Qualitative assessment of output embeddings

We visualize output representations of Llama 3 1B on addition and multiplication operations. For each expected result value y in the range 0-999,

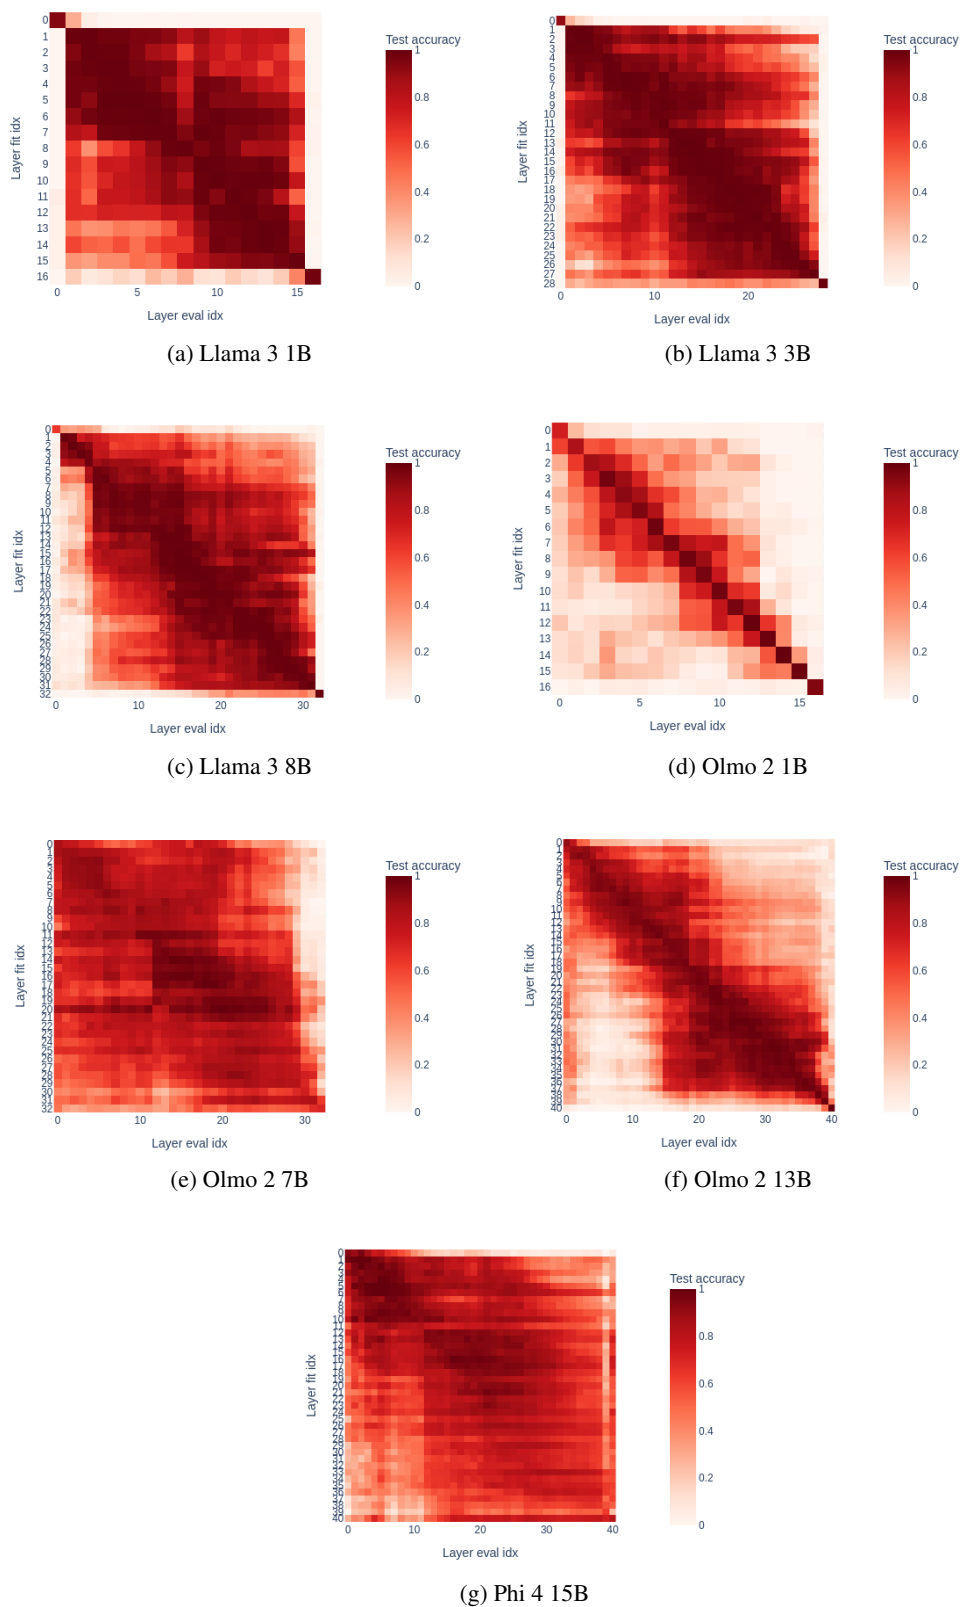
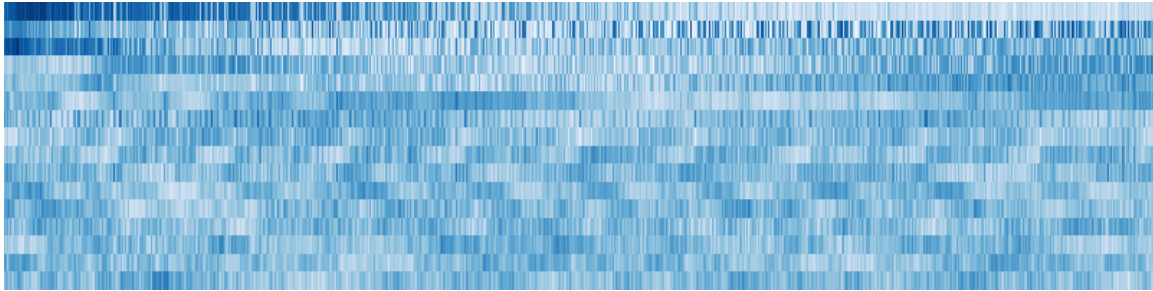
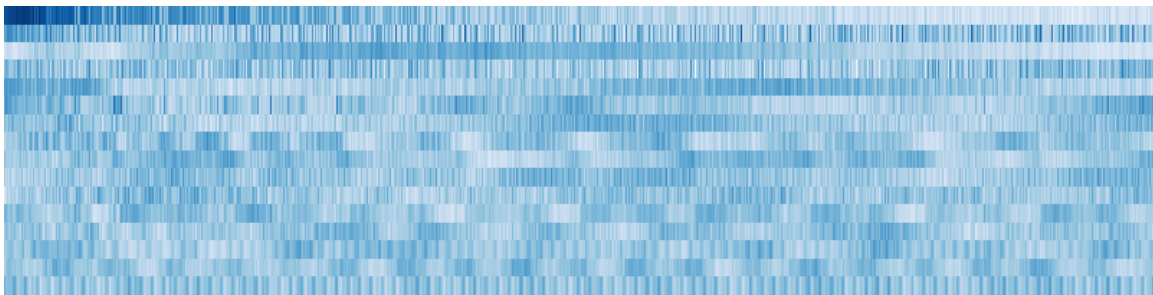


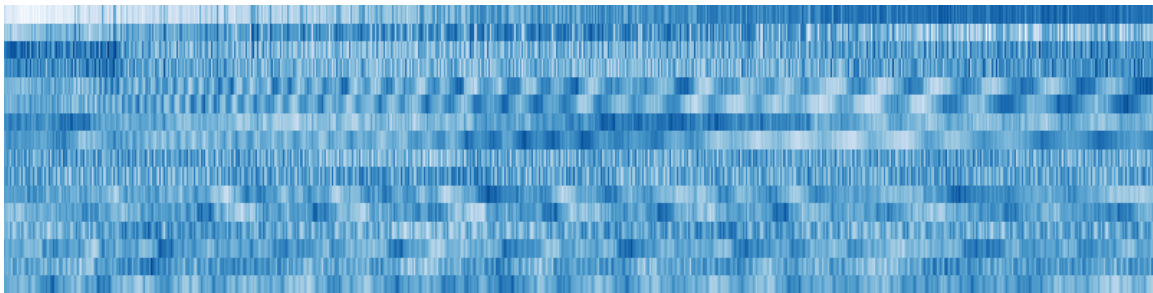
Figure 12: Probe fitted on one layer evaluated on all layers, synthetic math contexts. Olmo 2 1B shows the weakest cross-layer generalization among all models. Llama models display a strong separation between input/output embedding representations and the hidden representations.



(a) Llama 3 1B (layer 8/17)



(b) Olmo 2 7B (layer 16/31)



(c) Phi 4 15B (layer 20/41)

Figure 13: PCA of models' internal representations



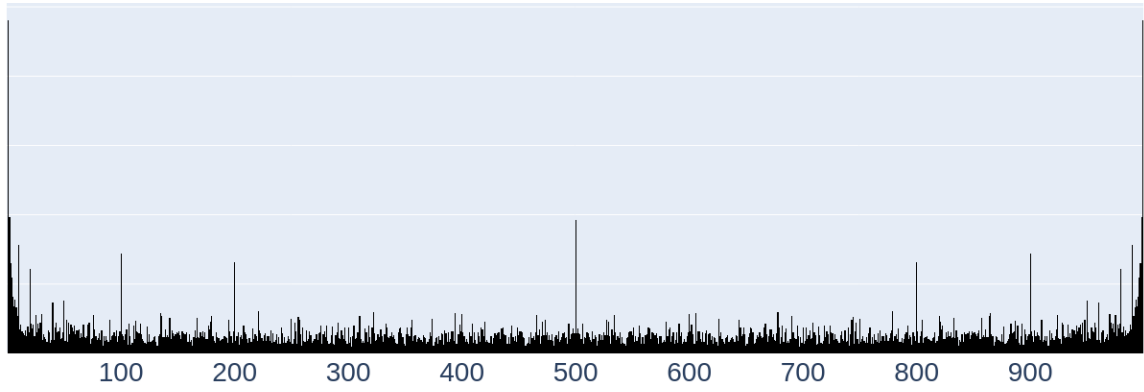


Figure 14: Maximal magnitudes of frequencies in Fourier transform of PCA of models' internal representations

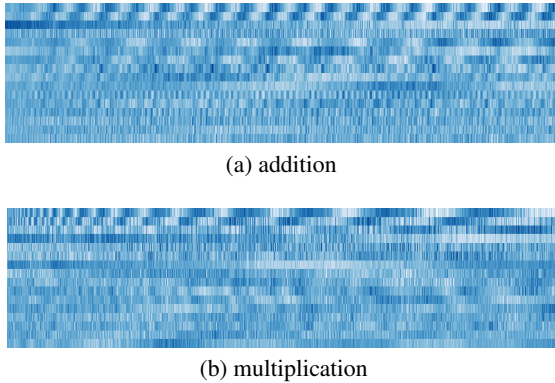


Figure 15: PCA of Llama 3 1B *output* representations for the operations of addition (recoverable with close-to 100% accuracy) and multiplication (recoverable with appx. 80% of accuracy)

we sample a random pair (x_1, x_2) from the same range, such that $x_1 + x_2 = y$ (or $x_1 \times x_2 = y$, respectively). We then predict the next token for prompts “ x_1 plus x_2 is ” and “ x_1 multiplied by x_2 is ” and collect the final output representations of the model before decoding. We then reduce the representations with PCA to 16 dimensions and visualize the result in Figure 15.

A.6 Probe sparsity

Cross-layer evaluation indicates that models’ internal representations tend to differ from the input/output embeddings. We find a notable discrepancy in the *sparsity* of their representation; Whereas the input/output embeddings represent numbers in a more *distributed* fashion, hidden layers use a small number of consistently-ordered sin features.

This is visualized in Figure 16, showing that intermediate layers in Llama 1B lead to probes with very few weights with values $> 10^{-5}$.

We hypothesize that this discrepancy can be caused by the models’ computational and optimization mechanics. As cross-entropy loss is minimized by confident one-hot predictions, models can benefit from using more features to create higher dissimilarity (separability) of tokens. In contrast, internal representations are not subject to direct optimization pressure and might benefit from fewer but more informative features.

Nevertheless, we note that this trend is pertinent across different models to a different extent: OLMo 7B and Phi 15B learn more distributed features also across the intermediate layers — which, in turn, also leads to a better generalization of universal

probes (Figure 5b), but still causing an outstanding drop in held-out accuracy compared to other layers. Finally, we note that the models’ sparsity profile does not relate to reported performances on arithmetic tasks (Kadlčık et al., 2025) and it also does not necessarily determine the accuracy of probes trained specifically for particular layer(s) (Figure 3).

A.7 Using probes for tracking representation consistency

We explore the origin of the high consistency of representations across layers within the internal mechanism that each transformer layer implements. In Figure 17, we visualize the accuracy of probing input representations from each of the components present in transformer layers (Llama 3.2 1B for brevity). We find the accurate numeric representation is scattered across different components in the attention mechanism, with the attention output projection largely violating the sinusoidal representation, which is then, to a large extent, reconstructed in the subsequent fully-connected block. Nevertheless, the consistency of layers’ output representation is primarily maintained by residual streams across layers.

A.8 Incremental refinement of model predictions on arithmetic tasks

We can use sinusoidal probes to estimate the absolute error per layer: the absolute difference between the predicted and target integer provides an estimate of how far off from the intended target a model is at any intermediate step of computation. Results aggregated per layer of Llama 3.2 3B are displayed in Figure 18. Across all operations, models tend to incrementally reduce the error towards the true answer value, *gradually* refining across layers. These results also suggest that there are particular layers responsible for an increase of errors in the model’s internal computation.

Figure 19 corroborates this hypothesis — showing that particular layers ‘break’ the correct result probed from the previous layer in large proportions of all predictions. For instance, in division, the 5th layer is responsible for breaking the correct result recoverable from the previous layer in over 10% of all cases. Further analyses, presented in Table 2, show that according to the probed representations, the models often achieve a correct result *internally*, even though it does *not* surface: this is the case of 56.8% of all surfaced (prediction) errors in sub-






Figure 16: **Difference in probe weights distribution (of W_{in}) across layers** shows that the initial and last layers represent numbers in a systematically sparser fashion.

Input	86	98	98	100	99	99	99	99	99	99	99	99	99	99	99	99	
Query	15	68	60	44	62	44	43	32	44	34	45	54	49	61	50	50	
Key	4	69	50	22	55	24	34	12	42	18	22	30	46	43	36	52	
Value	13	52	22	36	51	25	16	13	12	25	70	71	63	59	67	59	
Q+K+V	20	98	96	90	89	85	62	66	65	79	100	99	99	99	99	99	
Attn output	1	42	2	4	12	0	1	0	0	1	77	0	8	1	12	96	
MLP output	58	51	42	51	40	34	29	40	53	51	56	47	38	38	29	17	
+residuals	98	98	100	99	99	99	99	99	99	99	99	99	99	99	99	44	
		0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
		Layer															




(a) Addition (b) Subtraction


Figure 17: Accuracy of extracting input numbers from each component of the Transformer block using the sinusoidal probe.



(a) Addition (b) Subtraction



(c) Multiplication (d) Division



(c) Multiplication (d) Division

Figure 18: **Absolute error per layer:** Absolute error of numeric values probed from different layers.

traction, 26.3% in multiplication, and as much as 94.4% of all errors in division.

Having used our probes to identify key layers that are likely responsible for errors, we can now turn to practical **interventions** that they make possible. In practice, we simply remove the suspected

Figure 19: **Error aggregation per layer:** Relative ratio of cases where each of the layers in Llama 3.2 3B breaks the correct result from the previous layer.

layers from the model. This extremely naive intervention strategy is justified by the universality of representations across layers (Section 3.3). We remove one of the three layers with the highest error aggregation, separately in multiplication and in division (having a potential for improvement of accuracy). We find that this enhances performance in *four* out of six cases — in one case (layer 4) in multiplication, causing a 26% error reduction (from 90.38% to 92.91% and in *all* cases in division, with error reduction between 27–64%. However, we note that the absolute number of correct predictions probed in the lower layers is negligible, which hinders the general applicability of this methodology.

B Experimental details

The overview provided here is intended for informative purposes. For precise replication, we refer the reader to the companion codebase, which contains the exact hyperparameters for each experiment, available at <https://github.com/prompteus/numllama>.

In our probe architecture *param-sin*, we initialize amplitude parameters \mathbf{w}_a with a constant of 0.01, phase-shift parameters \mathbf{w}_ϕ randomly from uniform distribution between 0 and 2π , and frequency parameters \mathbf{w}_f with a vector of values regularly-spaced between 0 and $\frac{\pi}{4}$, and the parameters of linear projection \mathbf{W}_{in} randomly from uniform distribution between $-\sqrt{1/n}$ and $\sqrt{1/n}$, where d is the input dimension (number of model features).

We find that L1-regularization of the amplitude parameters \mathbf{w}_a improves generalization to held-out tokens by encouraging the probe to maintain small effective latent dimensionality even with a large actual dimension d . Similarly, we found similar benefits from regularization of the linear projection \mathbf{W}_{in} , although L2-regularization appears more robust. We find that an analogous configuration is also effective for the probe of Kadlčík et al., with L2-regularization of \mathbf{W}_{in} and L1-regularization of \mathbf{W}_{out} for inducing simplicity (sparsity) in the sinusoidal basis.

The probes are fitted on a train set, with a checkpoint selected on the validation set, and then evaluated on the held-out test set. In most experiments, we use the Adam optimizer with a learning rate of 10^{-4} and apply regularization with a coefficient of either 10^{-3} or 10^{-2} .

In multi-token decoding (Figure 6), we use a learning rate of $5 \cdot 10^{-4}$, 10,000 training steps, and early-stop training.

When fitting probes on natural language contexts (Figure 4, ‘Nat. lang. \rightarrow Nat. lang.’) and for cross-layer transfer (Figure 12), we use a hidden dimension of $d = 500$.

In our steering experiments (Figure 9), for f_{sin} , we optimize embeddings with respect to the probe using an SGD optimizer with a learning rate of 0.1, optimize the embeddings for 200,000 steps, and decrease the learning rate by a factor of 100 over the first 100,000 steps. As for $f_{\text{param sin}}$, we set regularization coefficient to 10^{-4} and use a hidden dimension of $d = 2000$.

We provide an overview of the datasets we use for natural language context experiments in Ta-

ble 3. In Section 4.1, we use BBC data as processed by Li et al. (2024), which we retrieved from [RealTimeData/bbc_news_alltime](https://github.com/RealTimeData/bbc_news_alltime).

In all cases, we trust the original creators of the datasets for handling private and sensitive information according to standard. As we do not redistribute the data, we take no specific steps towards verifying this point.

C Use of LLMs in writing and coding

LLM-based technology was used to polish writing and jump-start coding. Authors take full responsibility for the contents of this paper.

Domain	Dataset	Source	Numerical Context	Size
Culinary	Recipe NLG Lite	m3hrdadfi/recipe_nlg_lite	Quantities, measurements	6118
	FoodRecipe-ImageCaptioning	samsatp/FoodRecipe-ImageCaptioning	Ingredient amounts	719
Temporal	TimeLineExtraction	irlabamsterdam/TimeLineExtraction...CASE	Legal document dates	50
Arithmetic	MetaMathQA	meta-math/MetaMathQA	Mathematical reasoning	395K
	DROP	ucinlp/drop	Discrete reasoning	77.4K
	AQuA-RAT	deepmind/aqua_rat	Algebraic word problems	97.4K
Medical	ICD-10 Codes	atta00/icd10-codes	Diagnostic codes	25.7K
	ICD-10-CM	Gokul-waterlabs/ICD-10-CM	Medical classifications	74K

Table 3: Dataset specifications for numerical embedding analysis across natural language contexts.