

Encoder-only and decoder-only models understand abstract argument structure differently

Suet-Ying Lam^{1*†}, Qingcheng Zeng^{2*}, Rob Voigt³

¹UMass Amherst, ²Northwestern University, ³University of California, Davis

Abstract

Large language models (LLMs) based on encoder and decoder architectures exhibit strong performance across a range of tasks, yet systematic comparisons of how training objectives shape their behavior remain limited. Building on prior work establishing decoder-encoder differences (Weller et al., 2026), we ask whether this architectural difference shapes linguistic bias using two production-related tasks: semantic role assignment and syntactic structure choice (actives vs. passives). Across architectures, LLMs exhibit a robust animacy-first bias, preferring to place animate entities in subject position, diverging from human bias. However, when animacy cues are removed, encoders generalize human production strategy better than decoders and show a stronger preference for passive constructions. We argue that these differences reflect training objectives: encoders’ sensitivity to contextual compatibility supports flexible mapping from semantics to structure, whereas decoders’ reliance on surface frequency favors dominant patterns and disfavors infrequent structures such as passives.

1 Introduction

Encoder and decoder architectures differ fundamentally in their training objectives. Encoder-only models (henceforth encoders), such as BERT (Devlin et al., 2019), are typically trained with masked language modeling, predicting tokens from bidirectional context and encouraging representations that capture contextual compatibility. Decoder-only models (henceforth decoders), such as GPT-style models (Brown et al., 2020; OpenAI et al., 2024; Grattafori et al., 2024), are trained autoregressively to predict the next token given preceding context, directly modeling the probability of word sequences. Recent work using controlled comparisons between paired encoder and decoder models

trained on matched data and training recipes has shown that differences in training objectives alone lead to systematic differences in downstream behavior, such that encoders are better at classification and retrieval, whereas decoders are better at generation (Weller et al., 2026).

This paper asks whether these differences in training objectives also shape LLMs’ representations of verb semantics, and thus their performance on tasks that require verb semantics. We focus on two tasks that require knowledge of verb semantics, namely semantic role assignment and structural choice between actives and passives. Previous work that evaluated LLM performance on these two tasks mostly focused on BERT-variants (Petty et al., 2022; Papadimitriou et al., 2022; Ettinger, 2020; Wilson et al., 2023). Among those that included decoder-only models (Lee et al., 2024; Kauf et al., 2023), systematic comparisons controlling for training data, training recipe, and number of model parameters have not yet been conducted. It thus remains less clear if the reported sensitivity of LLMs in these works is more specifically a fact about encoder-only models, especially given the evidence that encoders and decoders encode semantic representations differently (Ethayarajh, 2019). Leveraging the paired encoder–decoder suite developed by Weller et al. (2026), this paper extends this line of work by targeting the source of LLMs’ biases, especially the role of training objectives, in these tasks.

Semantic roles refer to the roles of participants that involved in the event, e.g., who deliberately initiates and experiences the event (proto-agent) vs. who undergoes the event (proto-patient) (Dowty, 1991). When encoding an event from the conceptual level into a linguistic message, verb-specific constraints determine how semantic roles map onto grammatical functions. For instance, an action verb such as *chase* always places the agent-like role (the chaser) in the subject position in active construc-

*Equal contribution. †Corresponding author.

tions (e.g., *The dog* chased John) , while an emotion verb such as *amaze* always places the more agent-like experiencer to the object position in active constructions (e.g., The bird amazed *John*). Thus, the correct mapping between semantic role and grammatical function requires knowledge of verb semantics.

Verb semantics also play a crucial role in determining structural choice. Psycholinguistic research has found that human speakers tend to put the agent-like role to the subject position, and the patient-like role to the object position in actives (Ferreira, 1994; Do and Kaiser, 2022). Consequently, with emotion verbs that place the experiencer in object position, speakers are more likely to use passives to promote the experiencer to the subject position (e.g., *John was amazed by the bird*) than with action verbs that inherently align agents with subjects.

In this paper we present three major findings. First, LLMs exhibit a consistent animacy-first bias regardless of training objective. Second, when animacy cues are removed, encoders better generalize the human preference of using more passives with emotion verbs than action verbs, and larger models outperform smaller ones. This suggests that training objectives still matter for acquiring linguistic bias. Third, encoders show a substantially strong preference for passives compared to decoders, despite the infrequency of passives in human speakers. We suggest this is because the training objective of encoder-only models does not penalize the infrequency of passives: passives will not be rejected because they are less likely under the given context. Our codes and data are available [here](#).

2 Related Work

Since the introduction of the transformer architecture (Vaswani et al., 2017), substantial effort has been devoted to scaling encoder-only (Devlin et al., 2019; Liu et al., 2019; Warner et al., 2024), encoder-decoder (Raffel et al., 2020), and decoder-only

(Brown et al., 2020; OpenAI et al., 2024; Grattafiori et al., 2024) architectures into increasingly capable language models. Previous work has attempted to compare the difference between these architectures: For instance, Ethayarajh (2019) examined BERT and GPT-2 embeddings through semantic textual similarity tasks, showing that both architectures learn highly contextualized representations but encode them in different ways. Recently, Weller et al. (2026) pretrained a series of encoders and decoders under controlled head-to-head conditions and found that encoder models tend to perform strongly on classification and information retrieval tasks, whereas decoder-only models, often trained at much larger scales, exhibit strong generative capabilities. In this paper, we extend this line of work from task performance and representation analysis to language production, asking whether encoder and decoder architectures acquire different biases in thematic role assignment and syntactic structure choice. These tasks allow us to understand whether LLMs have an abstract representation of verbs such that they know the mapping between semantic and grammatical roles encoded in the verb (henceforth argument structure knowledge); or whether they rely on surface-level cues to assign semantic roles and choose between structures.

Prior work on LLM performance on semantic role assignment has produced mixed evidence: while some found that LLMs can distinguish plausible from implausible role assignments (Papadimitriou et al., 2022; Kauf et al., 2023; Ettinger, 2020), layer-wise analysis of BERT found no such sensitivity (Li et al., 2021). Furthermore, LLMs’ sensitivity to argument structure appears to depend heavily on surface-level cues available during training, breaking down when familiar verb representations are absent (Wilson et al., 2023) or when animacy contrasts between arguments are controlled for (Lee et al., 2024). However, these studies share a limitation beyond their inconsistent findings: they pre-

	<i>Animate noun in subject position</i>	<i>Inanimate noun in subject position</i>
Active	1. The activist protested/#angered the conflict.	2. The conflict #protested/angered the activist.
Passive	3. The activist was #protested/angered by the conflict.	4. The conflict was protested/#angered by the activist.

Table 1: Example frames constructed with the noun pair (*the activist* - *the conflict*), with the action verb (*protested*) or the emotion verb (*angered*). Frames that are unacceptable by human speakers are labelled with #.




Figure 1: Correctness of semantic role assignment by verb type (action vs. emotion), model type (encoder vs. decoder) and frame type (a: active frames vs. b: passive frames). In active frames, the correctness of semantic role assignment is higher with action verbs than with emotion verbs. In passive frames, the correctness is higher with emotion verbs than with action verbs.

dominantly test action verbs, where the animate entity canonically occupies the subject position in active sentences. It therefore remains unclear whether LLMs can correctly assign semantic roles in constructions that violate this pattern, such as emotion verbs, where the animate experiencer appears in the object position. Furthermore, none of these studies systematically examine whether model architecture shapes argument structure knowledge. We address both gaps by testing LLMs on action and emotion verbs across architecturally distinct models.

We then further extend the evaluation of argument structure knowledge beyond semantic role assignment to the choice of actives and passives, which vary by verb type and associated argument structure (Ferreira, 1994; Do and Kaiser, 2022). A model should demonstrate not only sensitivity to which noun fills which semantic role, but also knowledge of how those role-to-position mappings shift systematically across syntactic structures.

3 Experiments

3.1 Models

We evaluated a broad set of encoder-only and decoder-only language models to compare how architectural differences relate to the linguistic biases under investigation. Encoder models include ROBERTA-BASE/LARGE (Liu et al., 2019) and MODERNBERT-BASE/LARGE (Warner et al., 2024), covering both widely used BERT-style encoders and more recent encoder architectures. Decoder models include GPT2-SMALL/XL (Radford

et al., 2019), OLMo-2-7B-BASE (Team OLMo et al., 2025), and OLMo-3-7B-BASE (Team Olmo et al., 2026), covering earlier autoregressive models and recent open-weight decoder-only LLMs. In addition, to enable a controlled head-to-head comparison between architectures, we included the full Ettin series of paired encoders and decoders, ranging from 17M to 1B parameters (Weller et al., 2026). This model suite allows us to examine linguistic biases across architecture type, model family, scale, and training setup, while distinguishing broad encoder-decoder differences from patterns that may be specific to individual model families.

3.2 Materials

We constructed 88 different frames with the animate-inanimate noun pairs based on Ferreira (1994)’s experimental materials. Each frame comes with four versions, varying in noun order and sentence voice (active vs. passive). Each version is paired with both an action and an emotion verb, yielding $88 \times 8 = 704$ sentences. Table 1 illustrates all eight versions of sentences. This set of frames with animate-inanimate nouns were used for both the semantic role assignment task and the structural choice task. We further constructed 48 frames with two animate nouns (e.g., Ray stunned Max). These animate-only frames were used to test the possibility that any observed structural choice bias is due to animacy. Semantic role assignment was not evaluated for these frames because when both nouns are animate, there is no clear bias in assigning a semantic role (i.e., both *Ray stunned*

Max and *Max stunned Ray* are acceptable). All of the animate-only frames come with the eight versions listed in Table 1. The full set of frames can be found in Appendix E.

3.3 Evaluation

We evaluated the performance of encoder-only models by obtaining the masked probability of the verb token of each version of the animate-inanimate frame listed in Table 1, and the performance of decoder-only models by comparing the sentence probability of these frames.

For semantic role assignment, the masked probability of the verb or the sentence probability should be higher whenever the animate entity is placed in line with the verb-specific constraints than when it violates the verb-specific constraints. That is, the masked probability of the verb token or the sentence probability should be higher when the animate noun is the subject than when it is the object for action verbs (i.e., $1 > 2$; $4 > 3$ in Table 1), and when it is the object than when it is the subject for emotion verbs (i.e., $2 > 1$; $3 > 4$ in Table 1). For structural choice, we used the passive bias as the evaluation metric, defined as a higher masked probability of the verb or sentence probability of the correct passive frame than in the correct active frame (i.e., $4 > 1$ for action verbs; $3 > 2$ for emotion verbs). Recall that human speakers produce passives more frequently with emotion verbs than with action verbs (Ferreira, 1994). Thus, if semantic roles are correctly assigned, passive bias should be stronger for emotion verbs than for action verbs.

4 Results and Discussion

The results of the experiments are summarized in Figure 1 and 2. Figure 1 plots the correctness of semantic role assignment by model and verb type, in active and passive frames respectively. Figure 2 plots the passive bias of each model, i.e., the percentage of stimuli that passive frames are more probable than active frames, respectively with animate-inanimate and animate-only noun pairs.

Finding 1: Animacy-first bias across architectures. LLMs consistently place animate entities in subject position regardless of verb semantics, such that they make more mistakes in semantic role assignment for emotion verbs in active frames and action verbs in passive frames, where the verb only allows the animate experiencer or agent role to be put in the object position. This pattern holds across

both encoder and decoder architectures, suggesting it reflects a general property of LLM training rather than a specific architectural feature. For structural choice, when the noun pairs differ in animacy, all models exhibit a stronger passive bias in emotion verbs than in action verbs, in line with human results. When the animacy cue is removed, the passive bias only exists in a few models with a much smaller magnitude. This differs from human strategies, in which passive bias is found independent of the animacy effect (Ferreira, 1994).

These biases can be explained by an animacy-first heuristics, such that LLMs always prefer to put the animate entity at the first. This overgeneralization arises potentially because human language exhibits a pervasive preference for mentioning animate referents first, reflecting a broader cognitive bias observed not only in linguistic behavior (Leslie, 1995), but also across non-human species (New et al., 2007; Brocard et al., 2024).

Finding 2: Encoders generalize better and scaling matters. When animacy cues are removed for the structural choice task, larger models outperform smaller ones, and encoders are more likely to replicate the stronger passive bias in emotion verb than decoders overall.

Several factors may contribute to the generalization advantage of encoders. First, the bidirectional attention mechanism in encoders allows verb representations to be simultaneously informed by both the subject and object positions, potentially enabling encoders to develop more robust mappings between verbs and their thematic roles. Decoders, processing tokens left-to-right, may encode verbs in a way that is more strongly conditioned on the preceding subject, making them more susceptible to animacy cues at that position. Second, the masked language modeling objective encourages encoders to resolve semantic ambiguity using all available context, which may incidentally train them to rely more on verb-specific thematic role information rather than general animacy heuristics. Third, it is possible that encoder and decoder models develop systematically different verb embedding spaces, where encoder verb embeddings more clearly distinguish between verbs that strongly prefer animate subjects for semantic reasons and those that do not. This conjecture is partially confirmed by visualizing the embedding space using PCA and t-SNE projections of the largest paired encoders and decoders in the full Ettin series (Weller et al., 2026) in Appendix D.




Figure 2: Preference to use passive frames over active frames (passive bias) by verb type (action vs. emotion), model type and the animacy of noun pair (inanimate-animate vs. animate-only). Only frames for which the models assigned correct semantic roles are evaluated. In line with human bias, the passive bias is larger for emotion verbs than action verbs, but only when the noun pair differs in animacy.

Finding 3: Encoders prefer passives. Encoders exhibit a much stronger preference for passive constructions than decoders (>80%), regardless of the model size or the presence of animacy cues.

We attribute this to differences in training objectives: Decoders disfavor low-frequency constructions like passives, whereas masked probability generated by encoders essentially reflects the compatibility of the token with the bidirectional context. For a given verb, both actives and passives are appropriate contexts, so encoders are less likely to establish a bias towards one of the two structures. This account is supported by the representational analyses in Appendix D: encoders encode animacy/order and verb type as continuous semantic dimensions, whereas decoders encode these distinctions in a more discrete, locally clustered manner. Although neither architecture has a dimension for voice, this suggests that passive preference may not be directly encoded but instead emerges from how semantic factors are structured, given that encoders encode the semantic factors relevant to structural choice more systematically than decoders. This provides a plausible representational basis for their stronger passive preference.

5 Conclusion

This paper examined how encoder-only and decoder-only models differ in their argument structure knowledge of the mapping between semantic and grammatical role. Across the semantic role as-

signment task and the structural choice task, LLMs exhibit a robust animacy-first bias, diverging from human preferences, which further take the alignment between grammatical position and semantic role into account. However, the performance of the two architectures differ in their sensitivity to animacy: when animacy cues are absent, encoders generalize better but show a much stronger passive bias, suggesting a greater influence of verb semantics on structural choice. Decoders, by contrast, rely more on surface frequency and show a weaker sensitivity to verb semantics. These findings indicate that architectural differences shape not only performance but may also influence the linguistic knowledge acquired by models.

Limitations

First, the present tasks are primarily rooted in semantic understanding. Both semantic role assignment and structural choice require models to use verb semantics, animacy, and thematic-role information to evaluate sentence-frame compatibility. Our conclusions should therefore be interpreted as applying to semantic-role- and structure-related biases, rather than to linguistic bias in general.

Second, our model comparison is limited by the models included and how behavior is operationalized. Although the Ettin suite provides controlled encoder-decoder pairs, other models differ in training data, tokenization, scale, and model family. Also, as encoder and decoder probabilities are com-

puted with architecture-specific procedures, absolute effect sizes should be compared with caution.

References

- Sarah Brocard, Vanessa AD Wilson, Chloé Berton, Klaus Zuberbühler, and Balthasar Bickel. 2024. A universal preference for animate agents in hominids. *Iscience*, 27(6).
- 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. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*, Red Hook, NY, USA. Curran Associates Inc.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Monica L Do and Elsi Kaiser. 2022. Sentence formulation is easier when thematic and syntactic prominence align: evidence from psych verbs. *Language, Cognition and Neuroscience*, 37(5):648–670.
- David Dowty. 1991. Thematic proto-roles and argument selection. *language*, 67(3):547–619.
- Kawin Ethayarajh. 2019. [How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 55–65, Hong Kong, China. Association for Computational Linguistics.
- Allyson Ettinger. 2020. [What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models](#). *Transactions of the Association for Computational Linguistics*, 8:34–48.
- Fernanda Ferreira. 1994. Choice of passive voice is affected by verb type and animacy. *Journal of Memory and Language*, 33(6):715–736.
- 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.
- Carina Kauf, Anna A Ivanova, Giulia Rambelli, Emanuele Chersoni, Jingyuan Selena She, Zawad Chowdhury, Evelina Fedorenko, and Alessandro Lenci. 2023. Event knowledge in large language models: the gap between the impossible and the unlikely. *Cognitive Science*, 47(11):e13386.
- Eun-Kyoung Rosa Lee, Sathvik Nair, and Naomi Feldman. 2024. A psycholinguistic evaluation of language models’ sensitivity to argument roles. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3262–3274.
- Alan M. Leslie. 1995. A theory of agency. In Dan Sperber, David Premack, and Ann Jame Premack, editors, *Causal cognition: A multidisciplinary debate.*, pages 121–149. Clarendon Press/Oxford University Press, New York, NY/US.
- Bai Li, Zining Zhu, Guillaume Thomas, Yang Xu, and Frank Rudzicz. 2021. [How is BERT surprised? lay-erwise detection of linguistic anomalies](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4215–4228, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *Preprint*, arXiv:1907.11692.
- Joshua New, Leda Cosmides, and John Tooby. 2007. Category-specific attention for animals reflects ancestral priorities, not expertise. *Proceedings of the National Academy of Sciences*, 104(42):16598–16603.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Isabel Papadimitriou, Richard Futrell, and Kyle Mahowald. 2022. [When classifying grammatical role, BERT doesn’t care about word order... except when it matters](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 636–643, Dublin, Ireland. Association for Computational Linguistics.
- Jackson Petty, Michael Wilson, and Robert Frank. 2022. [Do language models learn position-role mappings?](#) *CoRR*, abs/2202.03611.

Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(1).

Team Olmo, :, Allyson Ettinger, Amanda Bertsch, Bailey Kuehl, David Graham, David Heineman, Dirk Groeneveld, Faeze Brahman, Finbarr Timbers, Hamish Ivison, Jacob Morrison, Jake Poznanski, Kyle Lo, Luca Soldaini, Matt Jordan, Mayee Chen, Michael Noukhovitch, Nathan Lambert, and 50 others. 2026. [Olmo 3](#). *Preprint*, arXiv:2512.13961.

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 Tafford, Taira Anderson, David Atkinson, Faeze Brahman, Christopher Clark, Pradeep Dasigi, Nouha Dziri, and 24 others. 2025. [olmo 2 furious](#). *Preprint*, arXiv:2501.00656.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.

Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Weller, Oskar Hallström, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, Nathan Cooper, Griffin Adams, Jeremy Howard, and Iacopo Poli. 2024. [Smarter, better, faster, longer: A modern bidirectional encoder for fast, memory efficient, and long context finetuning and inference](#). *Preprint*, arXiv:2412.13663.

Orion Weller, Kathryn Ricci, Marc Marone, Antoine Chaffin, Dawn Lawrie, and Benjamin Van Durme. 2026. [Seq vs seq: An open suite of paired encoders and decoders](#). *Preprint*, arXiv:2507.11412.

Michael Wilson, Jackson Petty, and Robert Frank. 2023. [How abstract is linguistic generalization in large language models? experiments with argument structure](#). *Transactions of the Association for Computational Linguistics*, 11:1377–1395.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

A Computational Experiment Details

All model evaluations were conducted with deterministic scoring rather than free-form generation, implemented with the Hugging Face Transformers library (Wolf et al., 2020). For encoder-only masked-language models, we used `AutoModelForMaskedLM` to compute the probability of the target verb in the masked verb position. When a target verb was split into multiple subword tokens, we inserted the corresponding number of mask tokens and computed the joint probability of the verb span by sequentially unmasking the subword tokens from left to right. For decoder-only models, we used `AutoModelForCausalLM` to compute full-sequence log probabilities for the active and passive sentence frames. Because this procedure evaluates the likelihood of fixed strings, no decoding hyperparameters such as temperature, top- p , or maximum generation length were used.

For representational analyses, we extracted contextual verb embeddings using `AutoModel`. The full sentence was provided to the model without masking, and the hidden states corresponding to the verb span were mean-pooled across all subword tokens. Unless otherwise specified, embeddings were extracted from the final hidden layer. For encoder models, the resulting verb representation is conditioned on both left and right context; for decoder models, the verb-position representation is conditioned only on the verb and its left context due to causal attention. All experiments were run on a single NVIDIA H100 GPU and required approximately one GPU hour in total.

B License Statement

Materials. Our experimental stimuli are based on the psycholinguistic materials of Ferreira (1994), which have been widely used to study thematic role assignment and active/passive structural choice. We adapt these materials for controlled model evaluation by constructing active and passive sentence frames with animate-inanimate and animate-only noun pairs. The resulting stimuli and evaluation scripts will be released for non-commercial academic research with attribution to the original source.

Models. We evaluate publicly released pretrained language models under their respective model licenses and terms of use. The encoder models include ROBERTA-BASE/LARGE (Liu et al.,

2019), MODERNBERT-BASE/LARGE (Warner et al., 2024), and the encoder members of the Ettin suite (Weller et al., 2026). The decoder models include GPT2-SMALL/XL (Radford et al., 2019), OLMo-2-7B-BASE (Team OLMo et al., 2025), OLMo-3-7B-BASE (Team Olmo et al., 2026), and the decoder members of the Ettin suite (Weller et al., 2026). All models are used only for non-commercial academic research.

C GenAI Statement

We disclose the use of generative AI tools in this work in accordance with the ACL Policy on the Use of AI Writing Assistance.

AI use in research artifacts. Generative AI tools were not used to create the experimental stimuli, compute model scores, extract embeddings, generate statistical results, or produce the figures. The stimuli were manually constructed from prior psycholinguistic materials, and all model evaluations and visualizations were conducted using the scripts described in this paper.

AI use in experiments. The language models evaluated in this work are themselves the subject of study. Their use is fully described in Section 3 and Appendix A. They were not used as auxiliary annotators or data-generation tools.

AI use in writing. We used AI assistants, including ChatGPT, for surface-level writing support, such as grammar correction, sentence-level rephrasing, LaTeX formatting suggestions, and code organization. All scientific claims, experimental design choices, analyses, interpretations, and conclusions were authored and verified by the human authors. AI assistants were not used to generate citations, statistical results, or factual claims reported in this paper.

Responsibility. The authors take full responsibility for the content of this paper, including any text or code that may have been drafted or edited with AI assistance.

D Plots of verb embeddings

Figure 3 shows a PCA projection of the representations for the stimuli in the largest paired encoders and decoders, which have 1B parameters. Figure 4 shows a t-SNE visualization of the same representations. In both active and passive conditions, representations are primarily separated along PC1,

with decoders forming a compact cluster on the left and encoders occupying a broader region on the right. Within the encoder space, points form a diagonal gradient running from the upper-left (lower PC1, higher PC2) to the lower-right (higher PC1, lower PC2). Along this axis, conditions vary systematically: configurations in which the animate entity appears earlier tend to lie toward one end of the gradient, while those in which the inanimate entity appears earlier are shifted toward the opposite end. This reflects a continuous displacement of points along both PC1 and PC2, rather than separation into discrete clusters. Verb type (AP = action vs. TE = emotion) induces smaller, approximately orthogonal or local shifts within this space, producing partial offsets among points with the same animacy/order configuration but different verbs. In contrast, decoder representations remain tightly clustered with minimal displacement along either axis. The overall geometry is highly similar across voice conditions. This pattern indicates that encoders encode animacy/order as a continuous, low-dimensional semantic direction in representation space, with verb type modulating this structure in a secondary way, consistent with their stronger generalization and more systematic structural biases, whereas decoders do not exhibit comparable linear organization of these distinctions.

In both plots, encoders and decoders are separated along the horizontal axis, with encoders on the right and decoders on the left. Within the decoder space, points form two vertically separated clusters, and conditions with different animacy/order configurations tend to group within the same cluster, indicating that animacy/order is reflected as local grouping structure rather than a continuous gradient. In contrast, encoder points are more diffusely distributed, and differences in animacy/order appear as gradual shifts in local neighborhoods rather than forming discrete clusters, with nearby points sharing similar configurations. The spatial organization remains largely unchanged between active and passive conditions. This suggests that animacy/order is encoded in a more clustered, locally separable way in decoders, but as a continuous, globally structured dimension in encoders.

PCA and t-SNE jointly indicate that, while both architectures exhibit an animacy-first bias, encoders encode animacy/order (with verb type as a secondary factor) as continuous, globally structured semantic dimensions, whereas decoders encode these distinctions in a discrete, locally clus-




Figure 3: PCA of sentence representations for encoder and decoder (1B) models under active and passive conditions. Representations are primarily separated by architecture along PC1, with decoders forming a compact cluster and encoders occupying a broader, structured space. Within the encoder space, conditions are organized along a diagonal gradient reflecting animacy/order, with verb type introducing smaller local shifts. The overall geometry is highly similar across voice conditions.




Figure 4: t-SNE visualization of sentence representations for encoder and decoder models (1B) under active and passive conditions. Encoders and decoders occupy distinct regions, with decoders forming two discrete clusters that group conditions by animacy/order and verb type, while encoder representations are more diffusely distributed with gradual local variation. The spatial organization remains largely consistent across voice conditions.

tered manner. The absence of a voice-related dimension suggests that passive preference is not explicitly represented, but instead emerges from how semantic information is structured. In this respect, the more integrated representational geometry of encoders supports flexible mapping from semantics to syntactic choice, leading to stronger passive bias and improved generalization in the absence of animacy cues, while decoders' less structured representations favor reliance on surface frequencies.

E Materials

Table 2 and 3 list the materials to construct the frames, with inanimate-animate noun pairs and animate-only noun pairs respectively. For inanimate-animate noun pairs, we constructed active frames in the format of [animate noun] + [emotion/action verb] + [inanimate noun] (corresponding to (1) in Table 1) and [animate noun] + [emotion/action verb] + [inanimate noun] (corresponding to (2) in Table 1). Passive frames were constructed in the same format except for the verb form (e.g., protested → was protested by).

For animate-only pair, the same frames were used except that the inanimate noun is replaced by [Noun 2] in Table 3.

Inanimate noun	Animate noun	Emotion verb	Action verb
conflict	activist	angered	protested
news	dictator	alarmed	suppressed
insult	fighter	provoked	bellowed
massage	athlete	relaxed	requested
affair	housewife	shamed	shunned
lullaby	baby	soothed	booed
cake	child	tempted	crumbled
evidence	detective	troubled	transformed
lightning	artist	frightened	painted
jokes	audience	entertained	applauded
hairdo	mother	irritated	complimented
thunder	children	terrified	dreaded
theft	judge	disturbed	abhorred
image	painter	haunted	sketched
layoffs	manager	worried	ordered
slayings	citizens	shocked	loathed
gift	mother	pleased	wrapped
sirens	police	alerted	startled
cave	boys	scared	explored
T.V.	student	distracted	ignored
map	explorers	guided	possessed
textbook	student	bored	xeroxed
cheesecake	dieter	enticed	devoured
praise	child	encouraged	awaited
tragedy	family	stunned	mourned
profits	executive	impressed	insured
defeat	troops	demoralized	disregarded
championship	team	excited	secured
crime	community	appalled	detested
fireworks	child	thrilled	feared
delay	customer	annoyed	despised
toy	child	amused	hurled
catcalls	referee	enraged	ignored
puzzle	genius	challenged	unraveled
play	reviewer	captivated	criticized
scandal	politician	disgraced	denounced
phenomenon	scientist	intrigued	analysed
pimple	teenager	embarrassed	camouflaged
belches	socialite	disgusted	disguised
rumors	queen	offended	overlooked

Table 2: Materials with inanimate-animate noun pair

Noun 1	Noun 2	Emotion verb	Action verb
constituent	politician	incensed	endorsed
correspondent	ambassador	embarrassed	authorized
pianist	conductor	provoked	denounced
publisher	novelist	scandalized	exploited
sheriff	cowboy	challenged	avoided
sitter	youngster	amused	appreciated
doctor	student	worried	approached
vampire	maiden	frightened	ignored
critic	actor	demoralized	disregarded
clergyman	atheist	offended	mistrusted
lawyer	judge	persuaded	seized
lifeguard	swimmer	shocked	loathed
trainer	boxer	thrilled	timed
supervisor	interpreter	pleased	praised
corporal	guerrilla	alerted	surrendered
administrator	undergraduate	discouraged	interrupted
superintendent	investigator	baffled	inspected
examiner	competitor	encouraged	criticized
grandfather	grandmother	upset	analyzed
hairstylist	socialite	delighted	resented
journalist	dictator	enraged	censored
rescuer	survivor	impressed	evaded
inhabitant	missionary	bored	blessed
performer	spectator	entertained	applauded
instructor	principal	annoyed	despised
teenager	murderer	terrified	resisted
thief	cop	frustrated	escorted
refugee	fisherman	scared	identified
therapist	immigrant	disturbed	heeded
hypnotist	dieter	relaxed	trusted
employer	apprentice	excited	rejected
camper	warden	concerned	halted
Linda	Rachel	distracted	scorned
Helen	Judy	soothed	weighed
Ruth	Ann	angered	disliked
Kathy	Christine	awed	idealized
Lucy	Susan	disappointed	misunderstood
Ray	Max	stunned	concealed
Stephen	Rose	intrigued	embraced
Herman	Michael	appalled	blamed
Phil	Tom	alarmed	excused
Truman	Nelson	perplexed	disobeyed
Emma	Bill	bothered	dreaded
Paula	Chris	amazed	controlled
Mark	Marjorie	tempted	scratched
Roger	Harry	convinced	admired
Dean	Pat	enticed	labeled
Albert	Barbara	astounded	neglected

Table 3: Materials with animate noun pairs