

## Clause-internal or clause-external? Testing Turkish reflexive binding in adapted versus chain-of-thought large language models

Sercan Karakaş\*

**Abstract.** This study evaluates whether state-of-the-art large language models capture the binding relations of Turkish reflexive pronouns. We construct a balanced evaluation set of 100 Turkish sentences that systematically pit local against non-local antecedents for the reflexives *kendi* and *kendisi*. We compare two contrasting systems: an OpenAI chain-of-thought model optimized for multi-step reasoning and Trendyol-LLM-7B-base-v0.1, a LLAMA-2–derived model extensively fine-tuned on Turkish data. Antecedent choice is assessed using a combined paradigm that integrates sentence-level perplexity with a forced-choice comparison between minimally differing continuations. Overall, Trendyol-LLM favors local bindings in approximately 70% of trials, exhibiting a robust locality bias consistent with a preference for structurally proximate antecedents. By contrast, the OpenAI model (o1 Mini) distributes its choices nearly evenly between local and long-distance readings, suggesting weaker or less consistent sensitivity to locality in this binding configuration. Therefore, these results reveal a marked contrast in binding behavior across the two systems and motivate closer analysis of how model architecture, training data, and inference-time reasoning strategies shape the representation of Turkish anaphoric dependencies.

**Keywords.** Turkish; reflexive binding; large language models; chain-of-thought; perplexity; psycholinguistics

**1. Introduction.** Recent advances in Large Language Models (LLMs) have yielded systems that achieve state-of-the-art performance across a diverse array of tasks, including open-domain question answering, text generation, summarization, and creative writing (e.g., Anil et al. 2023; Grattafiori et al. 2024; Gemma Team et al. 2025; OpenAI 2023; Ramos et al. 2026; Team et al. 2024). Beyond their well-documented linguistic versatility, these models appear capable of tackling sophisticated problem-solving challenges in domains such as software engineering and symbolic mathematics. Yet a central question remains unresolved: do contemporary LLMs truly “reason” in a human-like, logically compositional sense, or do they succeed mainly by exploiting vast training corpora to match surface-level statistical patterns? A growing body of evidence suggests the latter explanation still dominates. Diagnostic studies argue that LLMs often default to probabilistic pattern matching rather than rule-governed logical inference, revealing systematic blind spots whenever a task demands out-of-distribution generalization or explicit manipulation of abstract variables.

For this reason, the present paper pursues two complementary goals. First, we re-examine how effectively LLMs internalize fine-grained grammatical constraints, an issue frequently probed by comparing the probability distributions assigned by a model to minimal sentence pairs with those elicited from human acceptability judgments (Linzen et al. 2016; Lau et al. 2017; Chowdhury & Zamparelli 2018; Futrell et al. 2018; Gulordava et al. 2018; Tran et al. 2018; Wilcox et al.

---

\* I am grateful to Anastasia Giannakidou and Chris Kennedy for their valuable feedback and support. I also thank the Tu+10 audience and the anonymous reviewers for their helpful comments and suggestions. Any remaining errors are my own. Author: Sercan Karakaş, University of Chicago (skarakas@uchicago.edu).

2018; Hu et al. 2020; Marvin & Linzen 2018; Kim et al. 2024; Ide et al. 2025; Karakaş 2026). Notably, nearly all prior work centers on English, an abundantly resourced language whose syntactic patterns dominate many training corpora, leaving open how models fare on lower-resourced languages. Second, we ask whether recent algorithmic innovations designed to endow LLMs with explicit multi-step reasoning—exemplified by Chain of Thought style prompting and fine-tuning protocols—can measurably improve a model’s grasp of subtle grammatical phenomena (Wei et al. 2022; Yao et al. 2023a).

To that end, we focus on Turkish, a morphologically rich and comparatively under-represented language, and target one particularly informative construction: third-person reflexive anaphors. Turkish supplies two forms, *kendi* and *kendisi*, whose distribution is sensitive to both structural locality and discourse-pragmatic factors. By probing two contrasting systems—OpenAI o1 Mini, advertised as a “reasoning-oriented” model, and Trendyol-LLM-7B-base-v0.1, a strong general-purpose Turkish model—we quantify how probability assignments and categorical preferences align with patterns reported in psycholinguistic experiments on human speakers.

**2. Overview of reflexives in Turkish.** Turkish grammar distinguishes between two formally separate but etymologically related third-person reflexive pronouns—*kendi* and *kendisi*—whose distribution has long occupied a central place in the generative literature on binding and anaphoric reference (Özsoy 1983; Kornfilt 2000; Gürel 2003; Yakut 2015). While both items are glossed as ‘self’, they diverge morphosyntactically: *kendi* is morphologically bare and freely inflects for case (*kendini*, *kendine*, etc.), whereas *kendisi* contains the third-person possessive suffix *-si*, a fact that has been tied to emphatic and logophoric uses. At first glance, this morphological asymmetry suggests a syntactic split: traditional accounts classify *kendi* as a structurally constrained reflexive and *kendisi* as a long-distance or discourse-sensitive anaphor.

Recent experimental investigations, however, complicate this neat dichotomy. Judgement and processing studies by Yakut (2015), Gračanin-Yüksek et al. (2017), Bakay & Dillon (2022), and Oğuz & Kaiser (2023) demonstrate that, under certain discourse configurations—particularly those involving perspective-taking or shifted logophoric centers—the two forms converge in acceptability, allowing antecedents that are not clause-internal and thereby contravening the locality requirement suggested in Binding Principle A (Chomsky 1981). These findings raise the possibility that Turkish reflexivization is partly regulated by pragmatic or semantic dimensions (such as viewpoint or topicality) that override purely structural constraints. To ground the discussion, illustrative minimal pairs are provided in examples (1a) and (1b) below. In each pair, we juxtapose a structurally local reading (canonically licensed under Principle A) with a discourse-licensed non-local reading that native speakers nevertheless judge acceptable. This empirical tension makes Turkish reflexives an ideal test bed for probing how language models, and by extension, theories of grammar, integrate syntactic locality with higher-level interpretive factors.

(1a) Ali<sub>i</sub> Fatma-nın kendi-ni izle-diğ-in-i gör-dü.  
 Ali Fatma-GEN self-ACC watch-NMLZ-3SG-ACC see-PST  
 ‘Ali<sub>i</sub> saw that Fatma<sub>j</sub> watched self<sub>i/j</sub>.’

(1b) Ali<sub>i</sub> Fatma-nın kendisi-ni izle-diğ-in-i gör-dü.  
 Ali Fatma-GEN self-ACC watch-NMLZ-3SG-ACC see-PST  
 ‘Ali<sub>i</sub> saw that Fatma<sub>j</sub> watched self<sub>i/j</sub>.’

Traditional descriptions of Turkish binding treat *kendi* as the prototypically local reflexive:

in its canonical use, *kendi* must be bound within the smallest clause (or binding domain) that contains it, conforming to Binding Theory’s Principle A, which requires an anaphor to have a c-commanding antecedent in a sufficiently local structural domain (George & Kornfilt 1981; Enç 1989; Kornfilt 2000; Göksel & Kerslake 2005). In simple matrix clauses and many subordinate environments, this predicts tight clause-boundedness: a matrix subject cannot ordinarily serve as the antecedent for a reflexive embedded in a lower clause, nor can non-c-commanding discourse topics rescue the dependency.

The empirical record is more complicated, however. A parallel strand of work reports that *kendi*, specifically when interpreted with a third-person singular antecedent, sometimes permits reference to an argument located outside its minimal binding domain, yielding what is often termed a long-distance reading (Yakut 2015). These exceptions cluster in particular syntactic and discourse configurations: deeply embedded clausal complements of attitude predicates, narrative perspectives that privilege the matrix subject as the “center of consciousness,” and contexts where emphasis or contrast appears to license a wider search space for antecedent resolution. Although the precise conditions remain debated and judgments vary across speakers, dialects, and experimental methods, the recurring observation is that structural locality is not an absolute for *kendi* in third-person singular contexts.

In addition to the so-called clause-bound reflexive *kendi*, Turkish grammar furnishes a second reflexive/anaphoric element, *kendisi*, whose referential freedom markedly surpasses the strict locality constraints typically associated with Binding Principle A.

Morphologically, *kendisi* contains the third-person possessive suffix *-si*, a diachronic clue to its historical development from emphatic or logophoric constructions and to its gradual semantic bleaching into a more general anaphoric form. Syntactically, however, it behaves as a strikingly “liberated” anaphor: native-speaker judgments and experimental results (Enç 1989; Kornfilt 1997) demonstrate that *kendisi* (i) comfortably binds to clause-internal antecedents, replicating the patterns expected for canonical reflexivization, (ii) readily forms dependencies with long-distance antecedents that originate several clauses higher in the structure, and (iii) can even refer to extra-sentential antecedents, discourse-level referents introduced in the broader context but not overtly realized within the sentence containing the anaphor itself.

These properties position *kendisi* at the intersection of reflexivity, logophoricity, and discourse anaphora, challenging a strictly locality-based view of binding in Turkish. Its mixed profile also raises important questions for theoretical models that assume a uniform typology of anaphors, suggesting instead that Turkish distinguishes reflexive elements along both structural and discourse-pragmatic dimensions. In this sense, *kendisi* provides an empirical testing ground for evaluating how syntactic, semantic, and pragmatic factors jointly shape anaphoric interpretation.

**3. Large language models as psycholinguistic subjects.** Large-scale language models trained solely on raw, unannotated text have repeatedly surprised linguists by capturing non-local syntactic dependencies that generative theory argues are mediated by hierarchical structure rather than linear proximity. Corpus-based probes and carefully controlled minimal-pair tests show, for example, that current transformers license filler–gap relations, tracking a fronted *wh* phrase and its downstream gap, far above chance (Wilcox et al. 2018) and reliably enforce subject–verb number agreement across intervening material in a manner that mirrors human judgments (Linzen et al. 2016). These feats are more striking because the models receive no explicit gram-

matical supervision; the relevant constraints must be abstracted from surface statistics alone. Yet most celebrated successes involve constructions whose forms supply rich, redundant cues: overt agreement morphology, *wh*-fronting markers, or fixed lexical collocations. In these high-signal environments, massive training corpora can permit a model to exploit robust statistical regularities that stand in for deeper grammatical knowledge. Much less is known about how well the same systems cope with dependencies that are rare, structurally opaque, or heavily conditioned by discourse-pragmatic factors, which are situations where surface strings offer few, if any, diagnostic patterns. To probe this gap, a growing literature treats LLMs as surrogate psycholinguistic subjects. Researchers present the network with minimal sentence pairs that differ only at a critical word or phrase whose grammatical status determines acceptability, then ask whether the model assigns higher probability (lower surprisal) to the variant favored by native speakers. Applied to phenomena ranging from agreement to filler–gap licensing, this methodology has produced a mixed record (Linzen et al. 2016; Lau et al. 2017; Futrell et al. 2018; Gulordava et al. 2018; Tran et al. 2018; Wilcox et al. 2018): transformer models excel when the construction is frequent and overtly marked, but their performance deteriorates on low-frequency or syntactically opaque dependencies, casting doubt on whether apparent successes reflect genuine rule induction or sophisticated pattern matching. These findings force the field to confront a fundamental question: What standard should determine linguistic adequacy for artificial language models? From an engineering perspective, the bar is clear: models ought to approximate human behavior across all constructions, including those rare in training data. From a scientific vantage, however, the aim is to isolate which grammatical regularities are in principle learnable from sequential input, suggesting experimental designs that scaffold exposure and give models a fair chance to display latent competence. Contemporary large language models routinely produce prose that is grammatically well-formed, syntactically diverse, and semantically appropriate, suggesting that they have largely mastered formal linguistic competence, which is the rule-based knowledge underlying natural language structure (Mahowald et al. 2024; Piantadosi 2024). By contrast, their functional competence, which is the ability to deploy general knowledge and reasoning through language, remains contentious (Bender & Koller 2020; Marcus 2020; Mahowald et al. 2024).

Although LLMs excel at an impressive array of benchmarks from coherent multi-paragraph generation (Brown et al. 2020) to sentiment analysis and logical inference (Devlin et al. 2019; Liu et al. 2019; Radford et al. 2019), closed-book question answering, theory-of-mind tasks (Kosiński 2023; Trott et al. 2023) and selected commonsense-reasoning challenges (Zellers et al. 2018), closer scrutiny shows that their successes often hinge on shallow word-co-occurrence cues. When such surface regularities are neutralized, model performance frequently collapses (Ullman 2023; Xiang et al. 2024; Navarro Carranza 2025; Mancoridis et al. 2025; Karakaş & Şimşek 2026a), which is regarded as a fragility that contrasts with the relative robustness of human comprehension (Dominguez-Olmedo et al. 2025).

3.1. CHAIN-OF-THOUGHT REASONING. One influential response to the brittleness of “direct-answer” prompting is to modify the *inference-time* behavior of LLMs rather than their parameters. Chain-of-Thought (CoT) prompting asks the model to produce an explicit sequence of intermediate steps that bridges the question to an answer, effectively turning a single-shot prediction into a multi-step computation (Nye et al. 2021; Wei et al. 2022; Kojima et al. 2022). Beyond simple prompting, related methods (i) sample multiple reasoning traces and aggregate them by voting (*self-consistency*) (Wang et al. 2023), (ii) decompose problems into ordered subgoals (*least-*

*to-most prompting*) (Zhou et al. 2022), (iii) interleave reasoning with external actions such as tool use or information seeking (*ReAct*) (Yao et al. 2023b), and (iv) search explicitly over candidate reasoning paths using branching and backtracking (*Tree-of-Thought*) (Yao et al. 2023a). Closely related proposals further structure intermediate computations as executable programs (*program-of-thought prompting*) (Chen et al. 2023) or add reflection/self-critique loops to improve deliberation (Madaan et al. 2023; Shinn et al. 2023). Empirically, these strategies often improve accuracy on problems that require compositional arithmetic, multi-hop inference, or constraint satisfaction, and they offer a procedural account of how an output is obtained rather than only what output is produced.

At the same time, it remains unclear whether CoT-style scaffolding strengthens a model’s *linguistic* generalizations in the same way it can strengthen its *task* performance. Grammatical dependencies such as binding are not typically solved by explicit multi-step deduction; instead, they depend on structural locality, hierarchical relations, and discourse constraints that must be represented in the model’s underlying distribution over continuations. This makes Turkish reflexive binding a useful test case: if CoT primarily helps by supporting deliberate search and verification, it may have limited impact on reflexive antecedent preferences; if, however, CoT training induces more systematic sensitivity to structure, it could measurably sharpen locality effects. We therefore contrast a reasoning-oriented model with a Turkish-adapted base model and examine whether CoT correlates with more human-like binding choices. The basic CoT setup is illustrated in (2).

Suppose  $Q$  is the question you ask and  $A$  is one possible answer. Imagine there is an unobserved chain of thought  $C = (c_1, \dots, c_k)$  that leads from  $Q$  to  $A$ . Chain-of-Thought modeling says that to get the overall probability of  $A$  given  $Q$ , you “average” over possible reasoning chains  $C$ . Formally:

$$(2) \quad P(A | Q) = \sum_C P(A | Q, C) P(C | Q)$$

Conditional reasoning-chain probability:

$$(3) \quad P(C | Q) = \prod_{i=1}^k P(c_i | Q, c_{<i})$$

This factorization mirrors how large language models generate text: they predict each token (here, each reasoning step) given the accumulated context so far. Because the entire chain is scored—not just a single hop—the model can explore multiple candidate paths and assign higher probability to those whose later steps repair or refine earlier ones.

**Reasoning chains:**

$$(4) \quad \hat{A}^* = \arg \max_a \sum_{j=1}^N \mathbf{1} [\hat{A}_j = a]$$

Instead of trusting a single chain of thought, the model can generate  $N$  different reasoning chains; each chain produces its own answer. We then tally those answers and pick the one that appears most often.

On the other hand, Trendyol-LLM is an open-weight, bilingual Turkish–English language model developed by Trendyol (Trendyol 2024). It is a decoder-only, auto-regressive transformer based on a LLaMA-family architecture (Touvron et al. 2023) and fine-tuned using Low-Rank

Adaptation (Hu et al. 2021). As a base model, Trendyol-LLM is not instruction-tuned for dialogue; instead, it is trained to predict the next token given left context. In our experiments we interface with Trendyol-LLM in this standard language-modeling regime: the model processes each stimulus as a single left-to-right sequence and assigns probabilities to the two possible continuations. Thus, while the OpenAI model under study combines large-scale language modeling with an additional reasoning-oriented inference pipeline, Trendyol-LLM instantiates a more “classical” setup in which binding preferences must be expressed through the raw next-token distribution. This contrast allows us to ask whether a conventional, Turkish-adapted transformer already exhibits robust locality effects in reflexive binding, and whether reasoning-oriented inference yields qualitatively different behavior on the same materials.

## 4. Experiments.

4.1. THE LARGE LANGUAGE MODELS USED. Our experiments juxtapose o1 Mini (OpenAI 2024) with Trendyol-LLM-7B-base-v0.1 (Trendyol 2024), built on the LLaMA-2 backbone (Touvron et al. 2023), in order to tease apart the contributions of explicit reasoning supervision and Turkish-specific domain adaptation. The o-series models are trained with a reinforcement-learning objective that rewards internally coherent chains of thought. In practice, the learner receives credit not only for producing the correct terminal answer but also for generating intermediate reasoning steps that adhere to global logical consistency constraints. The other model starts from LLaMA-2-7B and undergoes Low-Rank Adaptation (LoRA) (Hu et al. 2021) on ~10 billion tokens dominated by Turkish and bilingual web text. Fine-tuning modifies projection matrices in self-attention plus the feed-forward blocks (rank = 64,  $\alpha$  = 128, dropout = 0.05).

4.2. THE EVALUATION METRICS. Large language models have become a workhorse for cognitive modeling because their token-wise probability distributions (and derived measures such as surprisal) can be mapped onto behavioral measures such as reading times and onto neural responses (EEG/fMRI/ECOG), and are also related to graded acceptability (Hale 2001; Levy 2008; Smith & Levy 2013; Frank et al. 2015; Lau et al. 2017; Goodkind & Bicknell 2018; Schrimpf et al. 2021; Caucheteux & King 2022; Škrjanec & Demberg 2026).

Following this line of work, we evaluate Trendyol-LLM using *sentence-level perplexity* (PPL), defined for a sentence  $W = (w_1, \dots, w_N)$  as:

$$(5) \quad \text{PPL}(W) = \exp \left( -\frac{1}{N} \sum_{i=1}^N \log P(w_i | w_{<i}) \right)$$

Sentence-level PPL measures average surprisal (processing cost) per token: lower PPL means the sentence is, on average, more predictable given its context. We illustrate the evaluation using the Turkish stimulus in (6), which mirrors the reflexive configuration introduced earlier.

- (6) Ali<sub>i</sub> Fatma<sub>j</sub>-nin kendi-ni izle-diğ-in-i gör-dü.  
 Ali Fatma-GEN self-ACC watch-NMLZ-3SG-ACC see-PST  
 ‘Ali<sub>i</sub> saw that Fatma<sub>j</sub> watched self<sub>i/j</sub>.’

For each item, we generate two minimal continuations that force distinct antecedents for the embedded reflexive. Concretely, we compare a continuation that forces a *matrix-subject* antecedent (7a) to a continuation that forces an *embedded-subject* antecedent (7b).

(7a) Ali izle-di.  
Ali watch-PST  
'Ali watched.'

(7b) Fatma izle-di.  
Fatma watch-PST  
'Fatma watched.'

Each continuation in (7a) and (7b) instantiates a distinct candidate antecedent for the embedded reflexive in (6). We quantify the model’s binding preference by computing sentence-level perplexity for each continuation: the continuation that attains lower PPL reveals the model’s favored antecedent. Consequently, if continuation (7b) (“Fatma watched.”) yields lower perplexity than continuation (7a) (“Ali watched.”), we infer that the model has resolved the reflexive to *Fatma*. In the case of the o1 model, which does not expose raw probability scores, we instead adopt a forced-choice evaluation, presenting both continuations verbatim and recording which antecedent the model selects.

4.3. TEST SET CONSTRUCTION. The test set consists of 100 carefully constructed Turkish sentences, evenly divided between those containing the reflexive *kendi* (50 items) and those containing the reflexive *kendisi* (50 items). In every sentence, the reflexive is embedded in a nominalized clause, ensuring that two antecedent readings are simultaneously licensed: a local antecedent residing inside the embedded clause and a non-local antecedent in the matrix clause. To manipulate only the reflexive marking while keeping all other lexical material constant, we created minimal pairs, thereby isolating the effect of reflexive type. We further counterbalanced reflexive forms across identical syntactic frames and verbs to eliminate potential frequency and lexical biases. By holding syntactic configuration constant and balancing reflexive type, this design controls for both item frequency and lexical familiarity. Consider the examples in (8).

(8a) Ali<sub>i</sub> Fatma<sub>j</sub>-nın kendi-ni izle-diğ-in-i gör-dü.  
Ali Fatma-GEN self-ACC watch-NMLZ-3SG-ACC see-PST  
'Ali<sub>i</sub> saw that Fatma<sub>j</sub> watched self<sub>i/j</sub>.'

(8b) Ali<sub>i</sub> Fatma<sub>j</sub>-nın kendisi-ni izle-diğ-in-i gör-dü.  
Ali Fatma-GEN self-ACC watch-NMLZ-3SG-ACC see-PST  
'Ali<sub>i</sub> saw that Fatma<sub>j</sub> watched self<sub>i/j</sub>.'

## 5. Results.

5.1. EVALUATION PROCEDURE. Recall that we evaluated both models on a 100-sentence test set, balanced along one dimension: reflexive type (50 sentences with the bare reflexive *kendi* and 50 with the reflexive *kendisi*). Each sentence embeds the reflexive within a nominalized complement, allowing us to generate two minimal continuations per item: one that forces the reflexive to co-refer with the embedded subject (local reading) and one that forces it to co-refer with the matrix subject (non-local reading). For the open-weight model, we exploited access to raw log-probabilities to compute sentence-level perplexity (PPL) for each continuation.

Crucially, we then reduced this gradient signal to a categorical decision by treating the continuation with lower PPL as the model’s preferred binding. Whenever the embedded-subject continuation yielded lower perplexity than the matrix-subject continuation, we inferred a local




Figure 1. Proportion of trials in which o1-mini favored the non-local vs. local antecedent, shown separately for *kendisi* and *kendi*.




Figure 2. Trendyol-LLM-7B shows a robust locality effect: the local antecedent yields lower PPL in most trials for both *kendisi* and *kendi*.

preference; conversely, if the matrix-subject continuation yielded lower perplexity, we inferred a non-local preference. This sentence-probability readout approach is standard in work that derives acceptability or preference judgments from language model probabilities (Lau et al. 2017). For the closed model, by contrast, continuous probability scores are not exposed, so we employed a forced-choice paradigm: both continuations were submitted verbatim, and the model was recorded as selecting one antecedent or the other as its binding decision. In both cases, therefore, the dependent measure is the same—whether the model favored the local or the non-local antecedent—while the evaluation procedure leverages the most informative interface available for each architecture and avoids conflating performance with superficial output patterns.

Figure 1 plots the proportion of trials in which OpenAI o1 Mini favored the non-local versus local antecedent in both conditions, showing near-chance performance. Figure 2 shows that Trendyol-LLM-7B, by contrast, exhibits a robust locality effect, with the local antecedent yielding lower PPL in both *kendisi* trials and *kendi* trials.

For Trendyol-LLM, we exploited access to raw token probabilities to compute sentence-level perplexity for each continuation. We then reduced this gradient signal to a categorical binding decision by treating the continuation with lower PPL as the model’s preferred antecedent. This probability-based readout is standard in work that derives acceptability or preference judgments from language model probabilities (Lau et al. 2017). For o1-mini, probability scores are

Model	Local choices ( $n = 100$ )	95% BCa CI for $\pi$	Binomial $p$	Inference
OpenAI o1 Mini	48	.38–.58	.764	Indistinguishable from chance
Trendyol-LLM-7B	71	.61–.80	< .001	Robust locality bias

Table 1. The Bootstrap Analysis. BCa = bias-corrected and accelerated bootstrap, 10,000 resamples (Efron 1987; Efron & Tibshirani 1993; DiCiccio & Efron 1996).

not available (OpenAI 2024), so we used the forced-choice paradigm.

5.2. STATISTICAL ANALYSIS. A two-tailed binomial test ( $H_0 = 0.5$  local preference) confirms that o1 Mini’s behavior is indistinguishable from chance, whereas Trendyol’s local bias is highly significant ( $p < .001$ ). A  $\chi^2$  comparison further shows the two models differ reliably in their antecedent choices. For each sentence the model’s preferred binding is defined as the continuation with lower sentence-level perplexity. Under the null hypothesis of no structural preference ( $\pi = .50$  for local antecedents), the number of local choices follows a Binomial( $N = 100, \pi = .50$ ) distribution.

Table 1 summarizes the bootstrap analysis of each model’s tendency to prefer local antecedents. OpenAI o1-mini selected the local antecedent on 48 of 100 trials (95% BCa CI = .38–.58), a rate that does not differ from the 50% chance baseline (two-tailed exact binomial test:  $p = .76$ ). By contrast, Trendyol-LLM-7B selected the local antecedent on 71 of 100 trials (95% BCa CI = .61–.80), indicating a robust locality bias (two-tailed exact binomial test:  $p = 3.2 \times 10^{-5}$ ). Confidence intervals were computed via the bias-corrected and accelerated (BCa) bootstrap with 10,000 resamples, confirming that only Trendyol-LLM-7B shows a statistically reliable preference for local binding. To directly compare the two models’ binding preferences, we constructed a  $2 \times 2$  contingency table crossing Model (o1-mini vs. Trendyol-LLM-7B) with Binding Choice (Local vs. Non-local): o1-mini (48 local, 52 non-local) and Trendyol-LLM-7B (71 local, 29 non-local). A chi-square test of independence shows a reliable association between model type and binding choice:

$$\chi^2(1, N = 200) = 10.98, \quad p < .001$$

$$\phi = \sqrt{\frac{\chi^2}{N}} = 0.23$$

Based on the bootstrap results, the  $\chi^2$  test of independence (1 degree of freedom;  $N = 200$ ) is significant ( $p = .0009$ ), confirming that the two models differ reliably in their reflexive-binding strategies. The corresponding effect size ( $\phi = 0.23$ ) indicates a small-to-moderate association, suggesting that approximately 5.5% of the variance in binding choice ( $\phi^2 \approx 0.055$ ) is attributable to model identity (and, by extension, differences in training/adaptation).

**6. Discussion and conclusion.** This study set out to discover whether language-specific adaptation or reasoning-oriented supervision is the more decisive factor in aligning large language models with context-sensitive binding behavior of Turkish third-person reflexive pronouns. To that end, we constructed a 100-item, lexically balanced test set in which the reflexives *kendi* and *kendisi* appear inside embedded clauses, forcing a competition between a structurally local antecedent (the embedded subject) and a clause-external antecedent (the matrix subject). Model preferences were read off from sentence-level perplexity, the exponentiated average negative log-likelihood, which serves as a cognitively interpretable proxy for incremental processing diffi-

culty and has long been used to relate human judgment gradients to language model probabilities (Linzen et al. 2016; Lau et al. 2017; Futrell et al. 2018). For the closed model, we used prompting and the forced-choice paradigm. We compared two contrasting systems: OpenAI o1 Mini, a multilingual model fine-tuned with Chain of Thought supervision that encourages multi-step reasoning (Wei et al. 2022; Wang et al. 2023), and Trendyol-LLM-7B, a LLaMA-2 derivative subjected to a LoRA fine-tune on roughly ten billion Turkish tokens drawn from newswire, literature, and e-commerce text. Psycholinguistic evidence shows that Turkish speakers accept both clause-internal and clause-external antecedents, exhibiting at most a graded preference for the embedded subject, an effect that weakens when the matrix subject is thematically demoted or the discourse perspective shifts (Bakay & Dillon 2022; Gračanin-Yüksek et al. 2017; Oğuz & Kaiser 2023). OpenAI o1 Mini does not show a strong locality bias in this paradigm, which is qualitatively compatible with reports of non-categorical preferences in humans, even though its balance may simply reflect the absence of strong Turkish priors. By contrast, Trendyol-LLM-7B showed an ‘over-local’ bias: in ~70% of cases, the continuation conditioned on the clause-internal antecedent had lower perplexity. The divergence is sizable ( $\phi \approx .23$ ) and unlikely to reflect mere noise: Trendyol is consistently more certain that the embedded subject is the antecedent, whereas o1 Mini distributes its confidence more evenly across local and long-distance readings. At the same time, the precise frequency and distribution of *kendi* and *kendisi* and of the discourse configurations that license long-distance binding in Trendyol-LLM’s training data are not publicly documented, so any account in terms of training-corpus bias must remain necessarily tentative. Against this background, two plausible, but contingent, factors may underlie Trendyol’s stance. First, insofar as the model’s fine-tuning corpus is reported to be dominated by relatively formal Turkish prose, it may underrepresent logophoric or perspectival contexts that facilitate long-distance binding, thereby strengthening structural locality while attenuating discourse-based cues. Second, if clause-internal antecedence is more frequent for these reflexives in the training distribution, the model may have internalized a default local preference that persists even in experimentally balanced materials. We therefore interpret Trendyol’s locality effect as robust within our paradigm, while acknowledging that a definitive attribution to frequency or discourse-distribution effects awaits access to, or detailed statistics from, the training corpus. Furthermore, although Trendyol trims its vocabulary, byte-pair tokenization still fragments agglutinative morphology, and morpheme boundaries often co-occur with embedded-clause boundaries; the model may therefore overgeneralize clause-internal regularities that are artifacts of segmentation (Bayram et al. 2025b; Karakaş & Şimşek 2026b). A key limitation is that Trendyol-LLM’s public documentation does not report corpus statistics for *kendi* and *kendisi* or for the discourse contexts that license long-distance binding. Thus, we cannot determine whether the observed locality bias reflects binding constraints or frequency-/exposure-driven preferences.

As a partial external baseline, Aydın’s (2019) corpus study (~450,000 words) reports that *kendi* ranks among the 100 most frequent words in Turkish, suggesting a substantial distributional advantage for *kendi*-type reflexives in naturalistic usage, even though these figures should be interpreted as reference-corpus estimates rather than measurements of the model’s actual training distribution (Aydın 2019). These observations cast Chain of Thought in a more circumspect light. CoT makes intermediate inference steps explicit but cannot conjure lexical-syntactic priors absent from the training signal; indeed, it can co-exist with either an over-strong locality effect (Trendyol) or a balanced distribution (o1 Mini) depending on the underlying statistics. Yet, as the model card does not provide token-level statistics for individual lexical items, we are unable to

report the training frequencies of *kendi* and *kendisi* for this specific model, and we leave a direct analysis of training-corpus effects to future work, ideally in collaboration with the model developers.

For low-resource languages, this means that reasoning scaffolds must be paired with corpora that faithfully capture both structural and discourse-pragmatic contingencies. Practical levers include morphology-aware tokenizers that respect agglutinative boundaries (Bayram et al. 2025b,a; Lian et al. 2025) and rigorously filtered corpora free from machine-translated noise (Cengiz et al. 2025). Once such foundations are in place, inference-time boosters, multi-round test-time thinking or chain-of-preference optimization, may yield additional gains, but the present findings warn that these techniques will not, by themselves, align models with the graded, context-dependent binding patterns exhibited by humans. Expanding the diagnostic set to discourse-rich stimuli, incorporating newer checkpoints such as OpenAI o4-mini (OpenAI 2025) and DeepSeek-R1 (DeepSeek-AI 2025), and testing phenomena where pragmatic cues override syntactic locality will be essential for deciding whether o1 Mini’s balanced profile signals latent grammatical competence or merely random indifference. In sum, genuine human-like linguistic behavior appears to require a synergistic strategy that marries deep-reasoning objectives to balanced, morphologically informed domain data, rather than privileging either ingredient in isolation.

## References

- Anil, Rohan, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri et al. 2023. PaLM 2 technical report. 10.48550/arXiv.2305.10403.
- Aydın, İlker. 2019. A corpus study on narrative texts in Turkish language. *International Journal of Progressive Education* 15(4). 200–214. 10.29329/ijpe.2019.203.15.
- Bakay, Özge & Brian Dillon. 2022. C-command effects in binding third-person reflexives and pronouns in Turkish. Poster presented at the 35th Annual Conference on Human Sentence Processing (HSP 35), University of California, Santa Cruz, March 24–26, 2022. Virtual conference hosted by UC Santa Cruz; in-person satellite held at the University of Maryland.
- Bayram, M. Ali, Ali Arda Fincan, Ahmet Semih Gümüő, Sercan Karakaő, Banu Diri, Savaő Yıldırım & Demircan Çelik. 2025a. Tokens with meaning: A hybrid tokenization approach for NLP. 10.48550/arXiv.2508.14292. <https://arxiv.org/abs/2508.14292>.
- Bayram, M. Ali, Ali Arda Fincan, Ahmet Semih Gümüő, Sercan Karakaő, Banu Diri & Savaő Yıldırım. 2025b. Tokenization standards and evaluation in natural language processing: A comparative analysis of large language models on Turkish. In *2025 33rd Signal Processing and Communications Applications Conference*, 1–4. 10.1109/SIU66497.2025.11112220.
- Bender, Emily M. & Alexander Koller. 2020. Climbing towards NLU: On meaning, form, and understanding in the age of data. In Dan Jurafsky, Joyce Chai, Natalie Schluter & Joel Tetreault (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5185–5198. Online: Association for Computational Linguistics. 10.18653/v1/2020.acl-main.463.
- Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan et al. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, vol. 33, 1877–1901.
- Caucheteux, Charlotte & Jean-Rémi King. 2022. Brains and algorithms partially converge in natural language processing. *Communications Biology* 5(1). 134. 10.1038/s42003-022-03036-1.
- Cengiz, Ayő Aysu, Ahmet Kaan Sever, Elif Ecem Ümütlü, Naime őeyma Erdem, Burak Aytan, Büőra Tufan, Abdullah Topraksoy, Esra Darıcı & Çaęrı Toraman. 2025. Evaluating the quality of benchmark datasets for low-resource languages: A case study on Turkish.
- Chen, Wenhui, Xueguang Ma, Xinyi Wang & William W. Cohen. 2023. Program of thoughts

- prompting: Disentangling computation from reasoning for numerical reasoning tasks. 10.48550/arXiv.2211.12588.
- Chomsky, Noam. 1981. *Lectures on government and binding*. Dordrecht: Foris.
- Chowdhury, Shammur Absar & Roberto Zamparelli. 2018. RNN simulations of grammaticality judgments on long-distance dependencies. In Emily M. Bender, Leon Derczynski & Pierre Isabelle (eds.), *Proceedings of the 27th International Conference on Computational Linguistics*, 133–144. Santa Fe, New Mexico, USA: Association for Computational Linguistics. <https://aclanthology.org/C18-1012/>.
- DeepSeek-AI. 2025. DeepSeek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning. 10.48550/arXiv.2501.12948.
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee & Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT 2019*, 4171–4186. Association for Computational Linguistics. 10.18653/v1/N19-1423.
- DiCiccio, Thomas J. & Bradley Efron. 1996. Bootstrap confidence intervals. *Statistical Science* 11(3). 189–228. 10.1214/ss/1032280214.
- Dominguez-Olmedo, Ricardo, Florian Eddie Dorner & Moritz Hardt. 2025. Training on the test task confounds evaluation and emergence. In *International Conference on Learning Representations (ICLR)*. 10.48550/arXiv.2407.07890.
- Efron, Bradley. 1987. Better bootstrap confidence intervals. *Journal of the American Statistical Association* 82(397). 171–185. 10.1080/01621459.1987.10478410.
- Efron, Bradley & Robert J. Tibshirani. 1993. *An introduction to the bootstrap*. New York: Chapman & Hall/CRC.
- Enç, Mürvet. 1989. Pronouns, licensing, and binding. *Natural Language & Linguistic Theory* 7(1). 51–92. 10.1007/BF00141347.
- Frank, Stefan L., Leun J. Otten, Giulia Galli & Gabriella Vigliocco. 2015. The ERP response to the amount of information conveyed by words in sentences. *Brain and Language* 140. 1–11. 10.1016/j.bandl.2014.10.006.
- Futrell, Richard, Ethan Wilcox, Takashi Morita & Roger Levy. 2018. RNNs as psycholinguistic subjects: Syntactic state and grammatical dependency. 10.48550/arXiv.1809.01329.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière et al. 2025. Gemma 3 technical report. 10.48550/arXiv.2503.19786.
- George, Leland M. & Jaklin Kornfilt. 1981. Finiteness and boundedness in Turkish. In Frank Heny (ed.), *Binding and Filtering*, 105–127. Cambridge, MA: MIT Press.
- Göksel, Aslı & Celia Kerslake. 2005. *Turkish: A comprehensive grammar*. London: Routledge. 10.4324/9780203340769.
- Goodkind, Adam & Klinton Bicknell. 2018. Predictive power of word surprisal for reading times is a linear function of language model quality. In Asad Sayeed, Cassandra Jacobs, Tal Linzen & Marten van Schijndel (eds.), *Proceedings of the 8th workshop on cognitive modeling and computational linguistics (CMCL 2018)*, 10–18. Salt Lake City, Utah: Association for Computational Linguistics. 10.18653/v1/W18-0102. <https://aclanthology.org/W18-0102/>.
- Gračanin-Yüksek, Martina, Sol Lago, D. F. Şafak, O. Demir & B. Kırkıcı. 2017. The interaction of contextual and syntactic information in the processing of Turkish anaphors. *Journal of Psycholinguistic Research* 46. 1397–1425. 10.1007/s10936-017-9502-2.
- Grattafiori, Aaron, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman et al. 2024. Llama 3 herd of models. 10.48550/arXiv.2407.21783.
- Gulordava, Kristina, Piotr Bojanowski, Edouard Grave, Tal Linzen & Marco Baroni. 2018. Colorless green recurrent networks dream hierarchically. 10.48550/arXiv.1803.11138.
- Gürel, Ayşe. 2003. Is the overt pronoun constraint universal? evidence from L2 Turkish. In

- Juana M. Liceras, Helmut Zobl & Helen Goodluck (eds.), *Proceedings of the 6th Generative Approaches to Second Language Acquisition Conference (GASLA 2002)*, 130–139. Somerville, MA: Cascadilla Proceedings Project.
- Hale, John. 2001. A probabilistic Earley parser as a psycholinguistic model. In *Proceedings of NAACL 2001*, 159–166. Association for Computational Linguistics.
- Hu, Edward J., Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang & Weizhu Chen. 2021. LoRA: Low-rank adaptation of large language models. 10.48550/arXiv.2106.09685.
- Hu, Jennifer, Jon Gauthier, Peng Qian, Ethan Wilcox & Roger Levy. 2020. A systematic assessment of syntactic generalization in neural language models. In Dan Jurafsky, Joyce Chai, Natalie Schluter & Joel Tetreault (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 1725–1744. Online: Association for Computational Linguistics. 10.18653/v1/2020.acl-main.158.
- Ide, Yusuke, Yuto Nishida, Justin Vasselli, Miyu Oba, Yusuke Sakai, Hidetaka Kamigaito & Taro Watanabe. 2025. How to make the most of LLMs’ grammatical knowledge for acceptability judgments. In Luis Chiruzzo, Alan Ritter & Lu Wang (eds.), *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 7416–7432. Albuquerque, New Mexico: Association for Computational Linguistics. <https://aclanthology.org/2025.naacl-long.380/>.
- Karakaş, Sercan. 2026. Plausibility as commonsense reasoning: Humans succeed, large language models do not. In *Proceedings of the 15th workshop on cognitive modeling and computational linguistics (CMCL)*. Proceedings.
- Karakaş, Sercan & Yusuf Şimşek. 2026a. Benchmarking source-sensitive reasoning in Turkish: Humans and LLMs under evidential trust manipulation. In *Proceedings of the 15th workshop on cognitive modeling and computational linguistics (CMCL)*. Proceedings.
- Karakaş, Sercan & Yusuf Şimşek. 2026b. From lemmas to dependencies: What signals drive light verbs classification? In Kemal Oflazer, Abdullatif Köksal & Onur Varol (eds.), *Proceedings of the Second Workshop Natural Language Processing for Turkic Languages (SIGTURK 2026)*, 220–227. Rabat, Morocco: Association for Computational Linguistics. 10.18653/v1/2026.sigturk-1.18.
- Kim, Huhng Joon, Youna Kim, Cheonbok Park, Junyeob Kim, Choonghyun Park, Kang Min Yoo, Sang goo Lee & Taeuk Kim. 2024. Aligning language models to explicitly handle ambiguity. 10.48550/arXiv.2404.11972.
- Kojima, Takeshi, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo & Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. In *Advances in Neural Information Processing Systems*, vol. 35, 22199–22213. 10.48550/arXiv.2205.11916.
- Kornfilt, Jaklin. 1997. *Turkish Descriptive Grammars*. London: Routledge.
- Kornfilt, Jaklin. 2000. Local and long-distance reflexives in Turkish. In Peter Cole, Gabriella Hermon & C.-T. James Huang (eds.), *Long Distance Reflexives*, vol. 33 Syntax and Semantics, 197–226. Brill. 10.1163/9781849508742\_007.
- Kosiński, Michał. 2023. Theory of mind may have spontaneously emerged in large language models. 10.48550/arXiv.2302.02083.
- Lau, Jey Han, Alexander Clark & Shalom Lappin. 2017. Grammaticality, acceptability, and probability: A probabilistic view of linguistic knowledge. *Cognitive Science* 41(5). 1202–1241. 10.1111/cogs.12414.
- Levy, Roger. 2008. Expectation-based syntactic comprehension. *Cognition* 106(3). 1126–1177. 10.1016/j.cognition.2007.05.006.
- Lian, Haoran, Yizhe Xiong, Zijia Lin, Jianwei Niu, Shasha Mo, Hui Chen, Peng Liu & Guiguang Ding. 2025. LBPE: Long-token-first tokenization to improve large language models. In *ICASSP 2025 - 2025 IEEE international conference on acoustics, speech and signal process-*

- ing (*ICASSP*), 1–5. 10.1109/ICASSP49660.2025.10889733.
- Linzen, Tal, Emmanuel Dupoux & Yoav Goldberg. 2016. Assessing the ability of LSTMs to learn syntax-sensitive dependencies. *Transactions of the Association for Computational Linguistics* 4. 521–535. 10.1162/tacl\_a\_00115.
- Liu, Yinhan, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer & Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. 10.48550/arXiv.1907.11692.
- Madaan, Aman, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon et al. 2023. Self-refine: Iterative refinement with self-feedback. In *Advances in Neural Information Processing Systems*, vol. 36. 10.52202/075280–2019.
- Mahowald, Kyle, Anna A. Ivanova, Idan A. Blank, Nancy Kanwisher, Joshua B. Tenenbaum & Evelina Fedorenko. 2024. Dissociating language and thought in large language models. *Trends in Cognitive Sciences* 28(4). 311–324. 10.1016/j.tics.2024.01.011.
- Mancoridis, Marina, Bec Weeks, Keyon Vafa & Sendhil Mullainathan. 2025. Potemkin understanding in large language models. 10.48550/arXiv.2506.21521.
- Marcus, Gary. 2020. The next decade in AI: Four steps towards robust artificial intelligence.
- Marvin, Rebecca & Tal Linzen. 2018. Targeted syntactic evaluation of language models. 10.18653/v1/D18–1151.
- Navarro Carranza, Juan Miguel. 2025. LLMs show surface-form brittleness under paraphrase stress tests. 10.48550/arXiv.2510.08616.
- Nye, Maxwell, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan et al. 2021. Show your work: Scratchpads for intermediate computation with language models. 10.48550/arXiv.2112.00114.
- Oğuz, Metehan & Elsi Kaiser. 2023. Interpretation of Turkish reflexive anaphors: Subjecthood vs. semantic role. In *Proceedings of the Workshop on Turkic and Languages in Contact with Turkic* 8, 158–172. 10.3765/55q87g40.
- OpenAI. 2023. GPT-4 technical report. 10.48550/arXiv.2303.08774.
- OpenAI. 2024. OpenAI o1 system card. <https://openai.com/index/openai-o1-system-card/>.
- OpenAI. 2025. OpenAI o3 and o4-mini system card. OpenAI System Card. Accessed 2026-02-03. <https://openai.com/index/o3-o4-mini-system-card/>.
- Özsoy, A. Sumru. 1983. *Kendi-reflexivization in Turkish: A syntactic, semantic and discourse analysis*: University of Michigan dissertation.
- Piantadosi, Steven T. 2024. Modern language models refute Chomsky’s approach to language. In Edward Gibson & Moshe Poliak (eds.), *From Fieldwork to Linguistic Theory: A Tribute to Dan Everett* (Empirically Oriented Theoretical Morphology and Syntax 15), 353–414. Berlin: Language Science Press. Chapter 15. Book DOI: 10.5281/zenodo.11351540.
- Radford, Alec, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei & Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog* 1(8). 1–9.
- Ramos, Miguel Moura et al. 2026. EuroLLM-22B: Technical report. 10.48550/arXiv.2602.05879.
- Schrimpf, Martin, Idan Asher Blank, Greta Tuckute, Carina Kauf, Eghbal A. Hosseini, Nancy Kanwisher, Joshua B. Tenenbaum & Evelina Fedorenko. 2021. The neural architecture of language: Integrative modeling converges on predictive processing. *Proceedings of the National Academy of Sciences* 118(45). e2105646118. 10.1073/pnas.2105646118.
- Shinn, Noah, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan & Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. In *Advances in neural information processing systems*, vol. 36. 10.52202/075280–0377.
- Škrjanec, Iza & Vera Demberg. 2026. Language models that match reader experience are better predictors of reading times. *Journal of Memory and Language* 146. 104677. 10.1016/j.

- jml.2025.104677.
- Smith, Nathaniel J. & Roger Levy. 2013. The effect of word predictability on reading time is logarithmic. *Cognition* 128(3). 302–319. 10.1016/j.cognition.2013.02.013.
- Team, Gemma, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot et al. 2024. Gemma 2: Improving open language models at a practical size. 10.48550/arXiv.2408.00118.
- Touvron, Hugo, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov et al. 2023. Llama 2: Open foundation and fine-tuned chat models. 10.48550/arXiv.2307.09288.
- Tran, Ke, Arianna Bisazza & Christof Monz. 2018. The importance of being recurrent for modeling hierarchical structure. In Ellen Riloff, David Chiang, Julia Hockenmaier & Jun’ichi Tsujii (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 4731–4736. Brussels, Belgium: Association for Computational Linguistics. 10.18653/v1/D18-1503.
- Trendyol. 2024. Trendyol-LLM-7b-base-v0.1. Hugging Face Model Card. Accessed 2024-10-29.
- Trott, Sean, Cameron Jones, Tyler Chang, James Michaelov & Benjamin Bergen. 2023. Do large language models know what humans know? *Cognitive Science* 47(7). e13309. 10.1111/cogs.13309.
- Ullman, Tomer. 2023. Large language models fail on trivial alterations to theory-of-mind tasks. 10.48550/arXiv.2302.08399.
- Wang, Xuezhi, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery & Denny Zhou. 2023. Self-consistency improves chain of thought reasoning in language models. 10.48550/arXiv.2203.11171.
- Wei, Jason, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi & Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, vol. 35, 24824–24837. 10.48550/arXiv.2201.11903.
- Wilcox, Ethan, Roger Levy, Takashi Morita & Richard Futrell. 2018. What do RNN language models learn about filler-gap dependencies? 10.18653/v1/W18-5423.
- Xiang, Zhen, Fengqing Jiang, Zidi Xiong, Bhaskar Ramasubramanian, Radha Poovendran & Bo Li. 2024. Badchain: Backdoor chain-of-thought prompting for large language models. 10.48550/arXiv.2401.12242.
- Yakut, Ayşe Büşra. 2015. The logophoric behavior of the strict local anaphor *kendi* ‘self’ in Turkish. In Andrew Joseph & Esra Predolac (eds.), *Proceedings of the 9th Workshop on Altaic Formal Linguistics (WAFL 9)* (MIT Working Papers in Linguistics 76). Cambridge, MA: MIT Working Papers in Linguistics.
- Yao, Shunyu, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao & Karthik Narasimhan. 2023a. Tree of thoughts: Deliberate problem solving with large language models. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt & S. Levine (eds.), *Advances in Neural Information Processing Systems*, vol. 36, 11809–11822. Curran Associates, Inc. 10.52202/075280-0517.
- Yao, Shunyu, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan & Yuan Cao. 2023b. ReAct: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations*. 10.48550/arXiv.2210.03629.
- Zellers, Rowan, Yonatan Bisk, Roy Schwartz & Yejin Choi. 2018. SWAG: A large-scale adversarial dataset for grounded commonsense inference. 10.18653/v1/D18-1009.
- Zhou, Denny, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans et al. 2022. Least-to-most prompting enables complex reasoning in large language models. 10.48550/arXiv.2205.10625.