

Dual Mechanism Priming Effects in Hindi Word Order

Sidharth Ranjan

IIT Delhi

sidharth.ranjan03@gmail.com

Marten van Schijndel

Cornell University

mv443@cornell.edu

Sumeet Agarwal

IIT Delhi

sumeet@iitd.ac.in

Rajakrishnan Rajkumar

IISER Bhopal

rajak@iiserb.ac.in

Abstract

Word order choices during sentence production can be primed by preceding sentences. In this work, we test the DUAL MECHANISM hypothesis that priming is driven by multiple different sources. Using a Hindi corpus of text productions, we model lexical priming with an n-gram cache model and we capture more abstract syntactic priming with an adaptive neural language model. We permute the preverbal constituents of corpus sentences, and then use a logistic regression model to predict which sentences actually occurred in the corpus against artificially generated meaning-equivalent variants. Our results indicate that lexical priming and lexically-independent syntactic priming affect complementary sets of verb classes. By showing that different priming influences are separable from one another, our results support the hypothesis that multiple different cognitive mechanisms underlie priming.

1 Introduction

Gries (2005) defines syntactic priming as the tendency of speakers “to repeat syntactic structures they have just encountered (produced or comprehended) before”. Starting with Bock (1986), a long line of experimental and corpus-based work has provided evidence for this phenomenon in the context of language production (see Reitter et al., 2011, for a through review). More recently, comprehension studies have also attested priming effects in a wide variety of languages (Arai et al., 2007; Tooley and Traxler, 2010), where prior experience of a syntactic structure alleviates the comprehension difficulty associated with subsequent similar syntactic structures during reading. The experimental record also demonstrates that lexical repetition affects syntactic priming (Reitter et al., 2011, and references therein). According to the DUAL MECHANISM ACCOUNT proposed by Tooley and Traxler

(2010), lexically independent syntactic priming effects are caused by an implicit learning mechanism (Bock and Griffin, 2000; Chang et al., 2006), whereas lexically dependent priming effects are caused by a more short-term mechanism, such as residual activation (Pickering and Branigan, 1998).

In the present work, we test this hypothesis of a dual mechanism of priming by analyzing whether different kinds of intersentential priming can account for the word order of different constructions in Hindi. Our main contribution is that we deploy precisely defined quantitative cognitive factors in our statistical models along with minimally paired alternative productions, whereas most previous experimental and corpus studies on priming only employ one or the other.

Hindi has a flexible word order, though SOV is the canonical order (Kachru, 2006). To investigate constituent ordering preferences, we generate meaning-equivalent grammatical variants of Hindi sentences by linearizing preverbal constituents of projective dependency trees of the Hindi-Urdu Treebank corpus (HUTB; Bhatt et al., 2009) of written text. We validated the assumptions underlying this method using crowd-sourced human judgments and compared the performance of our machine learning model with the choices made by human subjects. Pioneering studies of Hindi word order have demonstrated a wide variety of factors that influence order preferences, such as information status (Butt and King, 1996; Kidwai, 2000), prosody (Patil et al., 2008), and semantics (Perera and Srivastava, 2016; Mohanan and Mohanan, 1994). We incorporated measures of these baseline influences into a logistic regression model to distinguish the original reference sentences from our generated variants. We model lexical priming with an n-gram cache model and we capture more abstract syntactic priming with

an adaptive neural language model. Gries (2005) showed that syntactic priming effects are strongly contingent on verb class. To this end, we analyze model behavior on sentences involving the following verb classes: Levin’s (1993) syntactic-semantic verb classes, verbs involved in double object constructions, and conjunct verbs involving noun-verb complex predicates. To foreshadow our results, information-theoretic surprisal computed using our two different models predicts word order in complementary linguistic contexts over the baseline predictors. Moreover, for the task of choosing reference vs variant sentences, the model’s predicted choices matched the agreement between human subjects for all of Levin’s verb classes. By showing that different priming influences are separable from one another, our results support the dual mechanism hypothesis that multiple different cognitive mechanisms underlie priming.

2 Data

Our data set consists of 1996 reference sentences containing well-defined subject and object constituents corresponding to the projective dependency trees in HUTB corpus (Bhatt et al., 2009). The sentences in HUTB corpus belong to newswire domain and contains written text in naturally occurring context *i.e.*, every reference sentence in our dataset was taken from a newspaper article, thus situated in the context of preceding sentences. For each reference sentence in our data set, we created counterfactual grammatical variants expressing the same truth-conditional meaning¹ by permuting the preverbal constituents whose heads were linked to the root node in the dependency tree.² Inspired by grammar rules proposed in the NLG literature (Rajkumar and White, 2014), ungrammatical variants were automatically filtered out by detecting dependency relation sequences not attested in the original HUTB corpus. After filtering, we had 72833 variant sentences for our classification task.

¹A limitation of this definition: It does not capture the fact that, in contrast to marked orders, which necessitate context for a full interpretation, SOV canonical orders are neutral with respect to the preceding discourse (Gambhir, 1981).

²Appendix A explains our variant generation procedure in more detail.

3 Classification Task

In order to mitigate the data imbalance between the two groups (1996 references vs. 72833 variants), we follow Joachims (2002) by formulating our task as a pair-wise ranking problem.

$$w \cdot \phi(\textit{reference}) > w \cdot \phi(\textit{variant}) \quad (1)$$

$$w \cdot (\phi(\textit{reference}) - \phi(\textit{variant})) > 0 \quad (2)$$

The goal of the basic binary classifier model is shown in Equation 1, where the model learns a feature weight (w) such that the dot product of the variant feature vector ($\phi(\textit{variant})$) with w is less than the dot product of w with the reference feature vector ($\phi(\textit{reference})$). The same goal can be written as Equation 2 which ensures that w ’s dot product with the difference between the feature vectors is positive. This transformation alleviates issues from having dramatically unbalanced class distributions.

We first arranged the references and variants into ordered pairs (e.g., a reference with two variants would be paired as $(\textit{reference}, \textit{variant}_1)$ and $(\textit{variant}_2, \textit{reference})$), and then subtracted the feature vectors of the first member of the pair from the feature vectors of its second member. We then assigned binary labels to each pair, with *reference-variant* pairs coded as “1”, and *variant-reference* pairs coded as “0”, thus re-balancing our previously severely imbalanced classification task. Additionally, the feature values of sentences with varying lengths get centered using this technique. Refer to Rajkumar et al. (2016) and Ranjan et al. (2022b) for a more detailed illustration.

Using features extracted from the transformed dataset, we trained a logistic regression model to predict each reference sentence (see Equation 3). All the experiments were done with the Generalized Linear Model (GLM) package in *R*. Here *choice* is encoded by the binary dependent variable as discussed above (1: reference preference and 0: variant preference).

$$\textit{choice} \sim \begin{cases} \delta \text{ dependency length} + \\ \delta \text{ trigram surp} + \delta \text{ pcfg surp} + \\ \delta \text{ IS score} + \delta \text{ lexical repetition surp} + \\ \delta \text{ lstm surp} + \delta \text{ adaptive lstm surp} \end{cases} \quad (3)$$

3.1 Cognitive Theories and Measures

3.1.1 Surprisal Theory

According to the Surprisal Theory (Hale, 2001; Levy, 2008), comprehenders build probabilistic interpretations of phrases based on patterns they have already seen in sentence structures. Mathematically, the *surprisal* of the k^{th} word, w_k , is defined as the negative log probability of w_k given the preceding context:

$$S_k = -\log P(w_k|w_{1..k-1}) \quad (4)$$

These probabilities, which indicate the information load (or predictability) of w_k , can be calculated over word sequences or syntactic configurations. The theory is supported by a large number of empirical evidences from behavioural as well as broad-coverage corpus data comprising both comprehension (Demberg and Keller, 2008; Boston et al., 2008; Roark et al., 2009; Ranjan et al., 2022b; Staub, 2015; Agrawal et al., 2017) and production modalities (Demberg et al., 2012; Dammalapati et al., 2021, 2019; Ranjan et al., 2019, 2022a; Jain et al., 2018).

Using the above surprisal framework, we estimate various types of surprisal scores for each test sentence in our dataset as described below serving as independent variables in our experiment. The word-level surprisal of all the words in each sentence were summed to obtain sentence-level surprisal measures.

1. **Trigram surprisal:** We calculated the local predictability of each word in a sentence using a 3-gram language model (LM) trained on 1 million sentences of mixed genre from the EMILLE Hindi corpus (Baker et al., 2002) using the SRILM toolbox (Stolcke, 2002) with Good-Turing discounting.
2. **PCFG surprisal:** We estimated the syntactic probability of each word in the sentence using the Berkeley latent-variable PCFG parser³ (Petrov et al., 2006). We created 12000 phrase structure trees by converting HUTB dependency trees into constituency trees using the approach described in Yadav et al. (2017). Subsequently, we used them

³5-fold CV parser training and testing F1-score metrics were 90.82% and 84.95%, respectively.

to train the Berkeley PCFG parser. Sentence level log-likelihood of each test sentence was estimated by training a PCFG language model on four folds of the phrase structure trees and then testing on a fifth held-out fold.

3. **Lexical repetition surprisal:** Following the method proposed by Kuhn and De Mori (1990), we estimated cache-based surprisal of each word in a sentence using SRILM toolbox by interpolating a 3-gram LM with a unigram cache LM based on the history of words ($|H| = 100$) involving the preceding sentence with a default interpolation weight parameter ($\mu = 0.05$; see Equations 5 and 6). The basic idea is to keep track of word tokens that appeared recently and then amplify their likelihood of occurrence in the trigram word sequence. In other words, the following sentences are more likely to use words again that have recently appeared in the text (Kuhn and De Mori, 1990; Clarkson and Robinson, 1997). This way, we account for the lexical priming effect in sentence processing.

$$P(w_k|w_{1..k-1}) = \mu P_{\text{cache}}(w_k|w_{1..k-1}) + (1 - \mu) P_{\text{trigram}}(w_k|w_{k-2}, w_{k-1}) \quad (5)$$

$$P_{\text{cache}}(w_k|w_{1..k-1}) = \frac{\text{count}_H(w_k)}{|H|} \quad (6)$$

4. **LSTM surprisal:** The probabilities of each word in the sentence were estimated according to the entire sentence prefix using a long short-term memory language model (LSTM; Hochreiter and Schmidhuber, 1997) trained on 1 million sentences of the EMILLE Hindi corpus. We used the implementation provided in the neural complexity toolkit⁴ (van Schijndel and Linzen, 2018) with default hyperparameter settings to estimate surprisal using an unbounded neural context.
5. **Adaptive LSTM surprisal:** Following the method proposed by van Schijndel and Linzen (2018), we calculated the discourse-enhanced surprisal of each word in the sentence. The cited authors presented a simple way to continuously adapt a neural LM, and found that adaptive surprisal considerably outperforms

⁴<https://github.com/vansky/neural-complexity>

non-adaptive surprisal at predicting human reading times. They use a pre-trained LSTM LM and, after estimating surprisal for a test sentence, change the LM’s parameters based on the sentence’s cross-entropy loss. After that, the revised LM weights are used to predict the next test sentence. In our work, we estimated the surprisal scores for each test sentence using neural complexity toolkit by adapting our base (non-adaptive) LSTM LM to one preceding context sentence.

3.1.2 Dependency Locality Theory

Shorter dependencies are typically simpler to process than longer ones, according to the Dependency Locality Theory (Gibson, 2000), which has been demonstrated to be effective at predicting the comprehension difficulty of a sequence (Temperley, 2007; Futrell et al., 2015; Liu et al., 2017, cf. Demberg and Keller, 2008). Following the work by Temperley (2008) and Rajkumar et al. (2016), we calculated sentence-level dependency length by summing the head-dependent distances (measured as the number of intervening words) in the dependency trees of reference and variant sentences.

3.1.3 Information Status

Languages generally prefer to mention *given* referents, from earlier in the discourse, before introducing *new* ones (Clark and Haviland, 1977; Chafe, 1976; Kaiser and Trueswell, 2004). We assigned a *Given* tag to the subject and object constituents in a sentence if any content word within them was mentioned in the preceding sentence or if the head of the phrase was a pronoun. All other phrases were tagged as *New*. For each sentence, IS score was computed as follows: a) Given-New order = +1 b) New-Given order = -1 c) Given-Given and New-New = 0. For illustration, see Appendix B, which shows how givenness would be coded after a context sentence.

4 Experiments and Results

We tested the hypothesis that surprisal enhanced with inter-sentential discourse information (adaptive LSTM surprisal) predicts constituent ordering in Hindi over other baseline cognitive controls, including information status, dependency length, lexical repetition, and non-adaptive surprisal. For

our adaptation experiments, we used an adaptive learning rate of 2 as it minimized the perplexity of the validation data set (see Table 5 in Appendix C). The Pearson’s correlation coefficients between different predictors are displayed in Figure 2 in Appendix D. The adaptive LSTM surprisal has a high correlation with all other surprisal features and a low correlation with dependency length and information status score. On specific verbs of interest, we report the results of the regression and prediction experiments (using 10-fold cross-validation, i.e., a model trained on 9 folds was used to generate predictions on the remaining fold). A prediction experiment using feature ablation helped ascertain the impact of syntactic priming independent of lexical repetition effects. We conducted a fine-grained verb-specific analysis of priming patterns on conjunct verbs and Levin’s syntactic-semantic classes, followed by a targeted human evaluation of Levin’s verb classes.

4.1 Verb-Specific Priming

Individual verb biases are well known to influence structural choices during language production (Ferreira and Schotter, 2013; Thothathiri et al., 2017; Yi et al., 2019) and priming effects are also contingent on specific verbs (Gries, 2005). Therefore, we grouped Hindi verbs based on Levin’s syntactico-semantic classes using the heuristics proposed by Begum and Sharma (2017). Then we analyzed the efficacy of adaptive surprisal at classifying reference and variant instances of Levin’s verb classes (still training the classifier on the full training partition for each fold). Our results (Table 1, top block) indicate that the GIVE verb class was susceptible to priming, with adaptive surprisal producing a significant improvement of 0.12% in classification accuracy ($p = 0.01$ using McNemar’s two-tailed test) over the baseline model. The regression coefficients pertaining to Levin’s GIVE verb classes are presented in Table 6 in Appendix E. Other Levin verb frames did not show syntactic priming.

Our results align with previous work in the priming literature that shows GIVE to be especially susceptible to priming, thus providing cross-linguistic support to verb-based priming effects (Pickering and Branigan, 1998; Gries, 2005; Bock, 1986). The GIVE verb class in our data set includes different verbs that are semantically similar to *give* in En-

Type	Freq (%)	Baseline	Baseline + Adaptive LSTM
<i>Verb Class</i>			
DO	48.68	96.82	96.82
GIVE	19.35	93.86	93.98
SOCIAL	8.00	92.90	92.95
COMMUNICATE	6.25	93.94	93.98
LODGE	4.04	94.29	94.22
MOTION	3.87	90.87	90.76
PUT	2.97	95.28	95.28
DESTROY	2.42	95.58	95.63
PERCEPTION	0.73	87.48	87.10
OTHERS	3.69	90.63	90.22
<i>Alternations</i>			
S-DO	71.89	95.35	95.33
S-IO-DO	12.74	93.39	93.50
S-IO	15.37	94.98	95.04

Table 1: Prediction performance of verb-specific and subject-objects alternations (72833 points); Baseline denotes *base1* shown in Table 12; bold denotes McNemar’s two-tailed significance compared to baseline model in the same row)

glish, such as *de, saup, bhej, maang, dila, lautaa, vasul, thama, vaapas*. We found that all these verbs strongly exhibited double object constructions (Bergum and Sharma, 2017) and their arguments are often case marked (see Table 7 in Appendix F for more details).

4.2 Double Object construction

Previous studies on dative alternations in psycholinguistics have shown that the propensity of speakers to produce such constructions increases with their recent mention (Bock, 1986; Kaschak et al., 2006). The same factors also influence their predictability in reading comprehension (van Schijndel and Linzen, 2018; Tooley and Traxler, 2010; Tooley and Bock, 2014). To test whether such effects determine word-ordering decisions in Hindi, we isolated double object constructions from our dataset such that the main verb compulsorily has two objects *viz.*, direct and indirect objects in the sentence. Table 2 shows that all predictors (including adaptive and lexical repetition surprisal) are significant predictors of syntactic choice.

Then we analyzed the efficacy of adaptive surprisal at classifying reference and variant instances of double object constructions (still training the classifier on the full training partition for each fold). We also conducted a comparison of our results with single-object constructions. Our results (Table 1,

Predictor	$\hat{\beta}$	$\hat{\sigma}$	t
intercept	1.50	0.003	506.77
trigram surprisal	-0.14	0.017	-8.30
dependency length	0.02	0.003	6.20
pcfg surprisal	-0.11	0.005	-20.8
IS score	0.02	0.003	5.43
lex-rept surprisal	0.06	0.016	4.07
lstm surprisal	0.31	0.081	3.81
adaptive lstm surprisal	-0.59	0.081	-7.23

Table 2: Regression model on double object construction S-IO-DO data set (9278 data points; all significant predictors denoted by $|t|>2$)

bottom block) reveal that syntactic priming effects are present over and above lexical repetition effects. Syntactic priming is more influential in double object constructions (S-IO-DO) than in single object constructions (S-IO or S-DO), as attested by a significant improvement of 0.1% in classification accuracy ($p = 0.04$ using McNemar’s two-tailed test). Double object constructions are also highly case marked (see Table 8 in Appendix G) and 57.82% of these items contain verbs that belong to GIVE class (see Table 9 in Appendix H for more details). In the discussion section we present a more nuanced discussion on the effects of case-markers and a verb’s combinatorial properties on priming.

In summary, our analyses suggest that different verbs display varying strengths of priming effects, corroborating previous findings in the literature (Gries, 2005). Ditransitive constructions (denoted by S-IO-DO ordering) prime more strongly than other orderings, where verbs from the GIVE class strongly prefer canonical argument ordering⁵ while determining Hindi syntactic choices.

4.3 Example Analysis: Success of Adaptive LSTM Surprisal

We now discuss the example below to illustrate discourse-based syntactic priming effects (estimated via adaptive surprisal) in determining the preferred syntactic choice among referent-variant pairs (2a, 2b).

⁵For example, out of 284 instances, 89.79% of the lemma ‘de’ (GIVE class) occurs with canonical argument ordering in our test data set.

(1) **Context Sentence**

collingwood 8 aur jones 0 aur blackville 10-par
collingwood 8 and jones 0 and blackville 10-PSP
hi harbhajan-ki firki-ka sikaar **ban gaye**
EMPH harbhajan-GEN spin-GEN victim **become**-PST

*Collingwood became the victim of Harbhajan’s spin
on 8 and Jones on 0 and Blackville on just 10.*

- (2) a. *plunket* 14-par pathan-ki gend-par
plunket 14-PSP pathan-GEN ball-PSP
Gambhir-ko kaetch **de baethe**
gambhir-GEN catch give.PST.SG
(Reference)

*Plunket ended up giving a catch to Gambhir
on 14 off Pathan’s bowling.*

- b. 14-par *plunket* pathan-ki gend-par gambhir-ko
kaetch **de baethe** **(Variant)**

The LSTM LM when adapted to the previous sentence (1) and tested on referent-variant pairs (2) assigns a lower surprisal to the reference sentence (2a) than its competing variant (2b). It is conceivable that the adaptive LSTM surprisal learns syntactic patterns in the context sentence and prefers the reference sentence (over the variant) owing to the similarities between the reference and context sentences. Every other predictor aside from adaptive LSTM surprisal fails to predict the corpus reference sentence over the paired variant, in spite of the fact that the reference sentence has canonical ordering and the alternative variant has non-canonical ordering. This could be attributed to multiple factors. For example, dependency length would prefer the variant since the long-short sequence (14 *par-plunket*) in the variant minimizes its dependency length unlike the short-long sequence (*plunket-14 par*) in the reference sentence. Similarly, the intra-sentential surprisal models make the wrong choice while processing the sentences because they possibly get locally garden pathed due to the two consecutive proper nouns (NPs) *viz.*, *plunket* and *pathan* (referring to 2 distinct individuals in the real world as opposed to *plunket pathan* referring to a single individual). Table 10 and Figure 3 in Appendix I present the sentence-level predictor values of reference-variant pairs (Example 2) and their information profiles respectively illustrating these patterns.

Predictor	$\hat{\beta}$	$\hat{\sigma}$	t
intercept	1.50	0.001	1379.73
trigram surprisal	-0.09	0.005	-15.27
dependency length	0.01	0.001	7.82
pcfg surprisal	-0.07	0.002	-35.55
IS score	0.02	0.001	13.70
lex-rept surprisal	-0.02	0.005	-2.98
lstm surprisal	-0.14	0.016	-8.60
adaptive lstm surprisal	-0.12	0.016	-7.40

Table 3: Regression model on conjunct verb data set ($N = 51617$; all significant predictors denoted by $|t| > 2$)

4.4 Conjunct Verb Construction

In this section, we go beyond Levin’s verb class and study the effects of priming on sentences containing conjunct verbs. Hindi conjunct verbs are NOUN-VERB complex predicates (CP) in which a highly predictable verb follows a nominal leading to a non-compositional meaning (Butt, 1995; Mohanan, 1994; Husain et al., 2014). For example, the complex predicates, such as *khyaal rakhna* (‘care keep/put’; ‘to take care of’) with non-compositional meaning are associated with conjunct verb construction in our dataset (marked with the POF dependency relation label in the HUTB corpus) unlike the predicate *guitar rakhna* (‘guitar keep/put’; ‘to put down or keep a guitar’) that has compositional meaning.

In particular, we examined the impact of adaptive LSTM surprisal in predicting corpus reference sentences amidst the variants on the subset of the data consisting of conjunct verbs. Prior work in sentence comprehension has investigated the effects of expectation and locality in Hindi conjunct verb constructions (Husain et al., 2014; Ranjan et al., 2022b). The conjunct verb subset in our dataset contains 40.68% of reference sentences out of 1996, leading to 51,617 data points (referent-variant pairs) for our classification task.

Our regression results (Table 3) demonstrate that all the measures considered in our work are significant predictors of syntactic choice in Hindi. The negative regression coefficient of adaptive LSTM surprisal indicates that noun-verb predicate structures are more common in the context of similarly occurring noun-verb predicate structures, thus providing preliminary indication of potential priming effects. Further corpus analysis revealed that

35% of conjunct verb marked context sentences preceded reference sentences with conjunct verb phrases in our dataset. Adding adaptive LSTM surprisal into the regression model containing all other predictors significantly improved the fit ($\chi^2 = 187.27$; $p < 0.001$).

We now examine the relative performance of adaptive LSTM surprisal on conjunct verb constructions above and beyond every other feature in the classification model. We also conduct a feature ablation study to ascertain the impact of syntactic priming (adaptive LSTM surprisal) independent of lexical priming (lexical repetition surprisal) in determining syntactic choices in Hindi. We used the model trained over the entire training partition for each fold from the full dataset and then tested only on the conjunct-verb test partition. We found that even for conjunct verb constructions (rightmost column of Table 12 in Appendix J), adaptive LSTM surprisal induced a significant increase of 0.04% in prediction accuracy ($p = 0.04$ using McNemar’s two-tailed test) over a baseline comprised of all predictors but lexical repetition surprisal. Adaptive LSTM surprisal ceased to be a general predictor when lexical repetition surprisal was incorporated into the classification model. This result provides an evidence for a generalized *lexical boost effect* in Hindi, which operates over verb classes (conjunct verbs here) and not simply string-identical verbs, validating similar findings in English (Snider, 2009).

Additionally, Table 12 in Appendix J also presents the results of our classification experiment on the full dataset (72833 points). The findings discussed above for conjunct verb construction extend to full data as well. Besides, the feature ablation experiments on both full dataset and conjunct verb subset also suggest that when lexical repetition is taken into account there is weak tendency for the individual to repeat their own syntactic construction from preceding contextual sentence except for certain constructions as discussed in the preceding sections. Interestingly, similar findings have been reported for English dialogue corpora as well (Healey et al., 2014; Green and Sun, 2021). Future work needs to perform principled investigation on Hindi spoken data to understand the divergence and commonalities among written and verbal communication, and to make more substan-

tial claims about priming in language production.

4.5 Example Analysis: Success of Lexical Repetition Surprisal

This section discusses the following example where lexical repetition surprisal estimated using n -gram cache LM is the only predictor that makes the right choice by choosing the reference sentence 4a and every other measures predict the alternative variant 4b as their preferred syntactic choice.

(3) Context Sentence

jailon-ki jo haalat hai usme kisi
prisons-GEN such condition be in that any
kaedi-ka paagal ho jana maamuli baat hai
prisoner-GEN insane become-FUT minor thing

Such are the conditions of prisons, it is a minor thing for any prisoner to go insane.

- (4) a. varshon-tak mukadamen-ka intejaar
for years trial-GEN waiting
jailon-mein sadate een vichaaraadheen
prisons-LOC rotting these under-trial
kaidiyon-ko avasaad-mein jaane-ko
prisoners-ACC depression-LOC go-INF
vivash kar deti hai (**Reference**)
compel do.PRS.SG
Waiting for trial for years compels these undertrial prisoners rotting in jails to go into depression..
- b. jailon-mein sadate een vichaaraadheen
kaidiyon-ko avasaad-mein jaane-ko vivash
varshon-tak mukadamen-ka intejaar kar deti
hai (**Variant**)

Table 10 in Appendix I presents the sentence-level predictor values for referent-variant pairs (Example 4). For both sentences, the trigram cache LM assigns a high probability to the word ‘jailon’ (*prisons*) as the word is mentioned⁶ in the preceding context sentence (Example 3). However, at the sentence level, the cache LM allocates low surprisal score to the reference sentence (4a), thus predicting it to be a best choice than the variant sentence (4b). Altogether, this analysis indicates that lexical repetition surprisal accounts for the word’s preference to be in a syntactic configuration where the sequence is more probable, favoring the corpus reference sentence. We also argue that the long subject phrase (*varshon tak mukadamen ka intejaar*) in the reference sentence is hard to interpret

⁶In contrast, the examples discussed in Section 4.3 denoting syntactic priming do not have content word repetition across sentences.

Levin's verb Type (item count)	Agreement (%) human:corpus	Model (%) corpus	Model (%) human
DO (32)	84.38	65.63	68.75
SOCIAL (30)	86.67	70	76.67
GIVE (46)	86.96	67.39	67.39
COMMUNICATION (26)	100	92.31	92.31
MOTION (9)	77.78	66.67	66.67
PUT (8)	100	75	75
LODGE (8)	100	100	100
PERCEPTION (4)	100	100	100
DESTROY (2)	100	100	100
OTHERS (2)	100	100	100
Total (167)	89.92	74.85	76.65

Table 4: Targeted human evaluation — **Agreement human/corpus**: Percentages of times human judgement matches with corpus reference choice; **Model corpus**: Percentages of corpus choice correctly predicted by the classifier containing all the predictors; **Model human**: Percentages of human label correctly predicted by the classifier containing all the predictors

in isolation (*i.e.*, in the absence of the previous context sentence 3), potentially affecting the intra-sentential surprisal estimation that does not factor in the context information from the preceding sentence. Moreover, due to its long-short constituent and NEW-GIVEN orderings, additional factors like dependency length and IS score do not favor the reference sentence too.

4.6 Targeted Human Evaluation

We conducted a targeted human evaluation to validate our order-permutation analysis and to compare the choices made by the machine learning model with those of native speakers of Hindi. To this end, we designed a forced-choice task and collected sentence judgments from 12 Hindi native speakers for 167 randomly selected reference-variant pairs in our data set. Participants were first shown the context sentence, and were then asked to judge the most likely following sentence amongst the reference-variant pair. Each sentence was assigned a human label of “1” if more than 50% participants voted for it, or else “0”.

The stimuli containing reference and variant sentences belong to either of the orderings: *Canonical* or *Non-canonical*. Table 4 presents the results of our experiment. Overall, of 167 human-validated pairs, 89.92% of the reference sentences originally appearing in the HUTB corpus were also preferred by native speakers compared to the artificially generated variants expressing the very same proposition. Across all construction types, the full

model was better at predicting human preferences (76.65%) than it was at predicting the corpus reference sentences (74.85%). Furthermore, Pearson’s correlation between classifier predictions and human judgments was 0.534, and between classifier predictions and corpus labels was 0.497. Moreover, across all of the analyzed verb classes, the classifier using all measures was as good or better at predicting human choices than it was at discriminating reference from variant sentences, indicating a promising ability for these measures to reproduce human behavior. Further work is required to tease apart the relative contributions of the different predictors in modelling human choices.

5 Discussion

Written text is a consequence of language production and is often edited to facilitate comprehension for the readers. According to Levelt’s (1989) language production model, speakers evaluate their own utterances by comprehending their own speech and make necessary adjustments to an utterance via a self-monitoring loop. Therefore, we interpret our results in the light of the DUAL MECHANISM ACCOUNT (Tooley and Traxler, 2010) described earlier in the introduction. This account makes claims pertaining to both production and comprehension and Tooley and Bock (2014) demonstrates the parity of syntactic persistence across both phenomena. Our results indicate that the dual mechanism account can be extended to postulate a viable model of priming effects in Hindi word order. Constituent ordering choices demonstrate both lexically independent syntactic priming as well as lexically dependent effects. We discuss how these two effects are induced by distinct underlying mechanisms (as stated at the outset), *viz.*, implicit learning (Bock and Griffin, 2000; Chang et al., 2006), and residual activation (Pickering and Branigan, 1998) respectively.

Previous work suggests that lexical overlap between prime and target sentences enhances syntactic priming (Pickering and Branigan, 1998; Gries, 2005). We also show that certain verb classes are more susceptible to priming than others. Specifically, GIVE verbs selecting double objects are most prone to priming, a case demonstrated in English as well (Gries, 2005), thus providing cross-linguistic support for the finding. Hindi conjunct

verbs in prime sentences trigger subsequent target sentences with conjunct verbs, and preverbal word order patterns for Hindi conjunct verbs are influenced by the repetition of lexical cues mentioned in the previous sentence. These two findings lend credence to the idea in the literature that lexical boost effects are attested for heads (conjunct verbs in this case) as well as other non-head lexical items (Reitter et al., 2011). The explanation for such effects stems from the residual activation theory (Pickering and Branigan, 1998) where activated lemmas (linguistic category and combinatory nodes) in the prime utterance retain their activation for a short time. The residue of such activation is transferred to the target lemma. Reitter et al. (2011) proffer an alternative explanation for lexical boost via spreading activation mechanism posited by the ACT-R framework of cognition.

However, we observe syntactic priming independent of lexical effects over and above lexical repetition in double object constructions. Our verb-specific priming analyses indicate that prime sentences need not share the same main verb as the target sentence; instead, successive sentences may have a similar argument structure (subcategorization frame), which enforces a tendency to repeat canonical structures. Tooley and Traxler (2010) show that such effects are best explained by the implicit learning account (Bock and Griffin, 2000; Chang et al., 2006), where language users unconsciously acquire abstract routines over a period of time. In stark contrast to short-lived residual activation accounting for lexical boost effects, Bock and Griffin (2000) showed that lexically independent syntactic priming effects persisted even when 10 intervening structures occurred between prime and target utterances. The relationship between prediction (quantified using our surprisal measures) and learning is made explicit in the P-chain framework of Dell and Kittredge (2013) connecting production and comprehension. According to P-chain assumptions, prediction error leads to implicit learning, which in turn helps the prediction system to adapt to less common structures (like double object constructions), which are known to induce higher priming strengths compared to commonplace structures (Ferreira, 2003; Jaeger and Snider, 2007; Bernolet and Hartsuiker, 2010).

While our results demonstrate priming at the

level of verb classes, Husain and Yadav (2020) showed that the combinatory properties of the verb need not be the sole driver of priming in Hindi. In their self-paced reading experiments involving identical critical verbs in both prime and target sentences, they observed faster reading times only in the target condition where nominals were marked by a locative case marker (in contrast to accusative and ergative conditions). Language-specific properties like case markers and the relationship between Hindi production and comprehension processes needs to be investigated more thoroughly by extending our preliminary human evaluation (via a simple forced choice task) using more fine-grained measures like reading aloud and silent reading times as proposed by Ranjan et al. (2022a).

Overall, in line with the assumptions of the DUAL MECHANISM ACCOUNT, our main findings suggest that Hindi word order choices are influenced by both lexically independent syntactic priming effects as well as lexically dependent priming effects. Future inquiries need to explore controlled experiments to corroborate the psychological reality of our current results.

Acknowledgements

We would like to thank the first author's dissertation committee members: Dr Mausam and Dr Samar Husain, and the members of the Cornell's C.Psyd research group for their invaluable comments and feedback on this work. We thank Cornell's ethics board for their approval on human data collection for this project and Rupesh Pandey for his logistical help in collecting human judgement data associated with this work. We are also indebted to the anonymous reviewers of ACL 2022, ACL ARR 2021 and COMCO 2021 for their detailed and insightful feedback. Finally, the last two authors acknowledge extramural funding from the Cognitive Science Research Initiative, Department of Science and Technology, Government of India (grant no. DST/CSRI/2018/263).

References

Arpit Agrawal, Sumeet Agarwal, and Samar Husain. 2017. Role of expectation and working memory constraints in Hindi comprehension: An eyetracking corpus analysis. *Journal of Eye Movement Research*, 10(2).

- Manabu Arai, Roger PG Van Gompel, and Christoph Scheepers. 2007. Priming ditransitive structures in comprehension. *Cognitive psychology*, 54(3):218–250.
- Paul Baker, Andrew Hardie, Tony McEnery, Hamish Cunningham, and Robert Gaizauskas. 2002. Emille: a 67-million word corpus of indic languages: data collection, mark-up and harmonization. In *Proceedings of LREC 2002*, pages 819–827. Lancaster University.
- Rafiya Begum and Dipti Misra Sharma. 2017. Development and analysis of verb frame lexicon for hindi. *Linguistics and Literature Studies*, 5(1):1–22.
- Sarah Bernolet and Robert J. Hartsuiker. 2010. Does verb bias modulate syntactic priming? *Cognition*, 114(3):455 – 461.
- Rajesh Bhatt, Bhuvana Narasimhan, Martha Palmer, Owen Rambow, Dipti Misra Sharma, and Fei Xia. 2009. A multi-representational and multi-layered treebank for Hindi/urdu. In *Proceedings of the Third Linguistic Annotation Workshop, ACL-IJCNLP '09*, pages 186–189, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Bock and Z. Griffin. 2000. The persistence of structural priming: Transient activation or implicit learning. *Journal of Experimental Psychology*, 2(120):177–192.
- J.Kathryn Bock. 1986. Syntactic persistence in language production. *Cognitive Psychology*, 18(3):355 – 387.
- Marisa Ferrara Boston, John Hale, Reinhold Kliegl, Umesh Patil, and Shravan Vasishth. 2008. Parsing costs as predictors of reading difficulty: An evaluation using the potsdam sentence corpus. *Journal of Eye Movement Research*, 2(1).
- Miriam Butt. 1995. *The structure of complex predicates in Urdu*. Center for the Study of Language (CSLI).
- Miriam Butt and Tracy Holloway King. 1996. Structural topic and focus without movement. In Miriam Butt and Tracy Holloway King, editors, *Proceedings of the First LFG Conference*. CSLI Publications, Stanford.
- Wallace Chafe. 1976. Givenness, contrastiveness, definiteness, subjects, topics, and point of view. *Subject and topic*.
- Franklin Chang, Gary S. Dell, and Kathryn Bock. 2006. Becoming Syntactic. *Psychological Review*, 113(2):234–272.
- H. H. Clark and S. E. Haviland. 1977. Comprehension and the Given-New Contract. In R. O. Freedle, editor, *Discourse Production and Comprehension*, pages 1–40. Ablex Publishing, Hillsdale, N. J.
- P.R. Clarkson and A. J. Robinson. 1997. Language model adaptation using mixtures and an exponentially decaying cache. In *In Proceedings of ICASSP-97*, pages 799–802.
- Samvit Dammalapati, Rajakrishnan Rajkumar, and Sumeet Agarwal. 2019. Expectation and Locality Effects in the Prediction of Disfluent Fillers and Repairs in English Speech. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 103–109, Minneapolis, Minnesota. Association for Computational Linguistics.
- Samvit Dammalapati, Rajakrishnan Rajkumar, and Sumeet Agarwal. 2021. Effects of Duration, Locality, and Surprisal in Speech Disfluency Prediction in English Spontaneous Speech. In *Proceedings of the Society for Computation in Linguistics*, volume 4, page 10.
- Gary Dell and Audrey Kittredge. 2013. Prediction, production, priming, and implicit learning:: A framework for psycholinguistics. In Michael K. Tanenhaus Montserrat Sanz, Itziar Laka, editor, *Language Down the Garden Path: The Cognitive and Biological Basis of Linguistic Structures*, Oxford Studies in Biolinguistics. Oxford University Press.
- Vera Demberg and Frank Keller. 2008. Data from eye-tracking corpora as evidence for theories of syntactic processing complexity. *Cognition*, 109(2):193–210.
- Vera Demberg, Asad B. Sayeed, Philip J. Gorinski, and Nikolaos Engonopoulos. 2012. Syntactic surprisal affects spoken word duration in conversational contexts. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, EMNLP-CoNLL '12*, pages 356–367, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Victor S Ferreira. 2003. The processing basis of syntactic persistence: We repeat what we learn. *44th Annual Meeting of the Psychonomic Society*.
- Victor S Ferreira and Elizabeth R Schotter. 2013. Do verb bias effects on sentence production reflect sensitivity to comprehension or production factors? *Quarterly Journal of Experimental Psychology*, 66(8):1548–1571.
- Richard Futrell, Kyle Mahowald, and Edward Gibson. 2015. Large-scale evidence of dependency length minimization in 37 languages. *Proceedings of the National Academy of Sciences*, 112(33):10336–10341.
- Vijay Gambhir. 1981. *Syntactic restrictions and discourse functions of word order in standard Hindi*.

- Ph.D. thesis, University of Pennsylvania, Philadelphia.
- Edward Gibson. 2000. [Dependency locality theory: A distance-based theory of linguistic complexity](#). In Alec Marantz, Yasushi Miyashita, and Wayne O’Neil, editors, *Image, Language, brain: Papers from the First Mind Articulation Project Symposium*. MIT Press, Cambridge, MA.
- Daniel Gildea and T. Florian Jaeger. 2015. [Human languages order information efficiently](#). *CoRR*, abs/1510.02823.
- Clarence Green and He Sun. 2021. [Global estimates of syntactic alignment in adult and child utterances during interaction: Nlp estimates based on multiple corpora](#). *Language Sciences*, 85:101353.
- Stefan Th. Gries. 2005. [Syntactic priming: A corpus-based approach](#). *Journal of Psycholinguistic Research*, 34(4):365–399.
- John Hale. 2001. [A probabilistic Earley parser as a psycholinguistic model](#). In *Proceedings of the second meeting of the North American Chapter of the Association for Computational Linguistics on Language technologies*, NAACL ’01, pages 1–8, Pittsburgh, Pennsylvania. Association for Computational Linguistics.
- Patrick GT Healey, Matthew Purver, and Christine Howes. 2014. [Divergence in dialogue](#). *PloS one*, 9(6):e98598.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Samar Husain, Shravan Vasishth, and Narayanan Srinivasan. 2014. [Strong expectations cancel locality effects: Evidence from Hindi](#). *PLOS ONE*, 9(7):1–14.
- Samar Husain and Himanshu Yadav. 2020. [Target complexity modulates syntactic priming during comprehension](#). *Frontiers in Psychology*, 11:454.
- T. Florian Jaeger and Neal Snider. 2007. Implicit learning and syntactic persistence: Surprisal and cumulatvity. *University of Rochester Working Papers in the Language Sciences*, 3:26–44.
- Ayush Jain, Vishal Singh, Sidharth Ranjan, Rajkrishnan Rajkumar, and Sumeet Agarwal. 2018. [Uniform Information Density Effects on Syntactic Choice in Hindi](#). In *Proceedings of the Workshop on Linguistic Complexity and Natural Language Processing*, pages 38–48, Santa Fe, New-Mexico. Association for Computational Linguistics.
- Thorsten Joachims. 2002. [Optimizing search engines using clickthrough data](#). In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’02, pages 133–142, New York, NY, USA. ACM.
- Y. Kachru. 2006. *Hindi*. London Oriental and African language library. John Benjamins Publishing Company.
- Elsi Kaiser and John C Trueswell. 2004. The role of discourse context in the processing of a flexible word-order language. *Cognition*, 94(2):113–147.
- Michael P Kaschak, Renrick A Loney, and Kristin L Borreggine. 2006. Recent experience affects the strength of structural priming. *Cognition*, 99(3):B73–B82.
- Ayesha Kidwai. 2000. *XP-Adjunction in Universal Grammar: Scrambling and Binding in Hindi-Urdu: Scrambling and Binding in Hindi-Urdu*. Oxford studies in comparative syntax. Oxford University Press.
- Roland Kuhn and Renato De Mori. 1990. A cache-based natural language model for speech recognition. *IEEE transactions on pattern analysis and machine intelligence*, 12(6):570–583.
- Willem J. M. Levelt. 1989. *Speaking: From Intention to Articulation*. MIT Press.
- Beth Levin. 1993. *English verb classes and alternations: A preliminary investigation*. University of Chicago press.
- Roger Levy. 2008. [Expectation-based syntactic comprehension](#). *Cognition*, 106(3):1126 – 1177.
- Haitao Liu, Chunshan Xu, and Junying Liang. 2017. [Dependency distance: A new perspective on syntactic patterns in natural languages](#). *Physics of Life Reviews*, 21:171 – 193.
- K.P. Mohanan and Tara Mohanan. 1994. Issues in word order in south asian languages: Enriched phrase structure or multidimensionality? In Miriam Butt, Tracy Holloway King, and Gillian Ramchand, editors, *Theoretical perspectives on word order in South Asian languages*, pages 153–184. Center for the Study of Language and Information, Stanford, CA.
- Tara Mohanan. 1994. *Argument structure in Hindi*. Center for the Study of Language (CSLI).
- Umesh Patil, Gerrit Kentner, Anja Gollrad, Frank Kügler, Caroline Féry, and Shravan Vasishth. 2008. [Focus, word order and intonation in Hindi](#). *Journal of South Asian Linguistics*, 1(1):55–72.
- C. K. Perera and A. K. Srivastava. 2016. [Animacy-based accessibility and competition in relative clause production in Hindi and malayalam](#). *Journal of Psycholinguistic Research*, 45(4):915–930.

- Slav Petrov, Leon Barrett, Romain Thibaux, and Dan Klein. 2006. [Learning accurate, compact, and interpretable tree annotation](#). In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th Annual Meeting of the Association for Computational Linguistics, ACL-44*, pages 433–440, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Martin J. Pickering and Holly P. Branigan. 1998. [The Representation of Verbs: Evidence from Syntactic Priming in Language Production](#). *Journal of Memory and Language*, 39(4):633–651.
- Rajakrishnan Rajkumar, Marten van Schijndel, Michael White, and William Schuler. 2016. [Investigating locality effects and surprisal in written english syntactic choice phenomena](#). *Cognition*, 155:204–232.
- Rajakrishnan Rajkumar and Michael White. 2014. [Better surface realization through psycholinguistics](#). *Language and Linguistics Compass*, 8(10):428–448. ISSN: 1749-818X.
- Sidharth Ranjan, Sumeet Agarwal, and Rajakrishnan Rajkumar. 2019. [Surprisal and Interference Effects of Case Markers in Hindi Word Order](#). In *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*, pages 30–42, Minneapolis, Minnesota. Association for Computational Linguistics.
- Sidharth Ranjan, Rajakrishnan Rajkumar, and Sumeet Agarwal. 2022a. [Linguistic Complexity and Planning Effects on Word Duration in Hindi Read Aloud Speech](#). In *Proceedings of the Society for Computation in Linguistics (SCiL)*, 5:11.
- Sidharth Ranjan, Rajakrishnan Rajkumar, and Sumeet Agarwal. 2022b. [Locality and expectation effects in hindi preverbal constituent ordering](#). *Cognition*, 223:104959.
- David Reitter, Frank Keller, and Johanna D. Moore. 2011. [A computational cognitive model of syntactic priming](#). *Cognitive Science*, 35(4):587–637.
- Brian Roark, Asaf Bachrach, Carlos Cardenas, and Christophe Pallier. 2009. [Deriving lexical and syntactic expectation-based measures for psycholinguistic modeling via incremental top-down parsing](#). In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1 - Volume 1, EMNLP '09*, pages 324–333, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Neal Snider. 2009. [Similarity and structural priming](#). In *Proceedings of the 31th Annual Conference of the Cognitive Science Society*, pages 815–820.
- Adrian Staub. 2015. [The effect of lexical predictability on eye movements in reading: Critical review and theoretical interpretation](#). *Language and Linguistics Compass*, 9(8):311–327.
- Andreas Stolcke. 2002. SRILM — An extensible language modeling toolkit. In *Proc. ICSLP-02*.
- David Temperley. 2007. [Minimization of dependency length in written English](#). *Cognition*, 105(2):300–333.
- David Temperley. 2008. [Dependency-length minimization in natural and artificial languages](#). *Journal of Quantitative Linguistics*, 15(3):256–282.
- Malathi Thothathiri, Daniel G. Evans, and Sonali Poudel. 2017. [Verb bias and verb-specific competition effects on sentence production](#). *PLOS ONE*, 12(7):1–18.
- Kristen M Tooley and Kathryn Bock. 2014. [On the parity of structural persistence in language production and comprehension](#). *Cognition*, 132(2):101–136.
- Kristen M Tooley and Matthew J Traxler. 2010. [Syntactic priming effects in comprehension: A critical review](#). *Language and Linguistics Compass*, 4(10):925–937.
- Marten van Schijndel and Tal Linzen. 2018. [A neural model of adaptation in reading](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4704–4710.
- Himanshu Yadav, Ashwini Vaidya, and Samar Husain. 2017. [Keeping it simple: Generating phrase structure trees from a Hindi dependency treebank](#). In *TLT*.
- Eunkyung Yi, Jean-Pierre Koenig, and Douglas Roland. 2019. [Semantic similarity to high-frequency verbs affects syntactic frame selection](#). *Cognitive Linguistics*, 30(3):601–628.

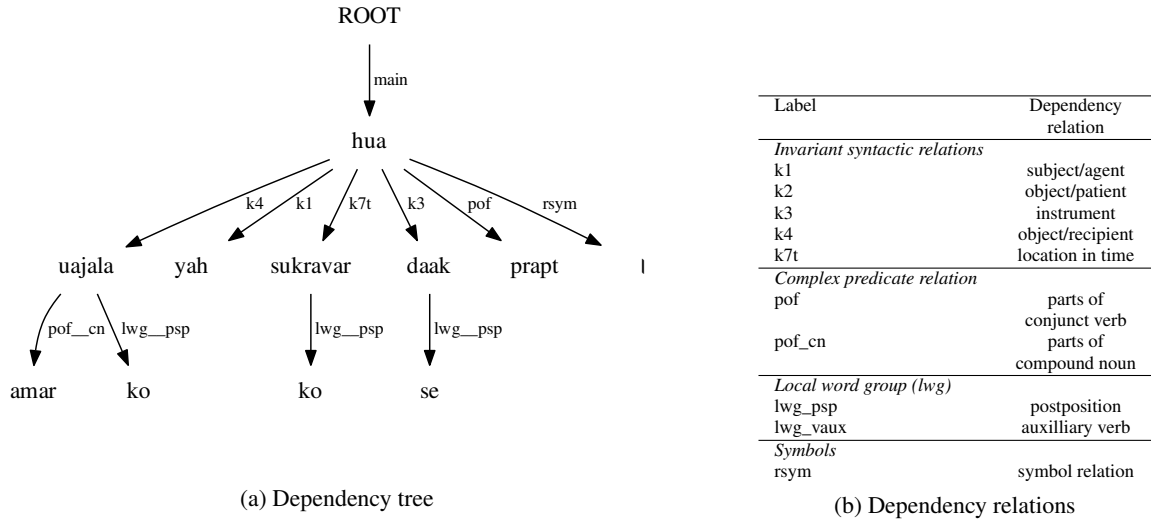


Figure 1: Example HUTB dependency tree and relation labels

Appendix

A Variant Generation

(5) Context sentence

amar ujala-ki bhumika nispaksh rehti hai
 Amar Ujala-GEN role unbiased remain be.PRS.SG

Amar Ujala’s role remains unbiased.

- (6) a. amar ujala-ko **yah** sukravar-ko daak-se prapt hua [Given-Given = 0] (**Reference**)
 Amar Ujala-ACC **it** friday-on post-INST receive be.PST.SG
 Amar Ujala received **it** by post on *Friday*.
- b. **yah** amar ujala-ko sukravar-ko daak-se prapt hua [Given-Given = 0] (**Variant 1**)
- c. sukravar-ko **yah** amar ujala-ko daak-se prapt hua [New-Given = -1] (**Variant 2**)

This work uses sentences from the Hindi-Urdu Treebank (HUTB) corpus of dependency trees (Bhatt et al., 2009) containing well-defined subject and object constituents. Figure 1 displays the dependency tree (and a glossary of relation labels) for reference sentence 6a. The grammatical variants were created using an algorithm that took as input the dependency tree corresponding to each HUTB reference sentence. The re-ordering algorithm permuted the preverbal dependents of the root verb and linearized the resulting tree to obtain variant sentences. For example, corresponding to the reference sentence 6a and its root verb “hai” (see figure 1a), the preverbal constituents⁷ with parents as “ujala”, “yah”, “suravar”, “daak”, and “prapt” were permuted to generate the artificial variants (6b and 6c). The ungrammatical variants were automatically filtered out using dependency relation sequences (denoting grammar rules) attested in the gold standard corpus of HUTB trees. In the dependency tree 1a, “k4-k1”, “k7t-k1”, “k3-k7t”, and

⁷Hindi is not a strictly verb-final language but the majority of the constituents in the HUTB corpus are preverbal. Ranjan et al. (2022b) in their corpus analysis with 13274 HUTB sentences found 20,750 pairs of preverbal constituents and 2599 pairs of postverbal constituents. Therefore, we also limit our variant generation (via reordering of constituents) and subsequent experiments on word-order variation in the preverbal domain only and leave the postverbal constituents in the reference-variants sentences as it is.

“pof-k3” are dependency relation sequences. In cases where the total number of variants exceeded 100 (a random cutoff),⁸ we chose 99 non-reference variants randomly along with the reference sentence.

B Information Status Annotation

The subject and object constituents in a sentence were assigned a *Given* tag if any content word within them was mentioned in the preceding sentence or if the head of the phrase was a pronoun. All other phrases were tagged as *New*. The sentence example 6 illustrates the proposed annotation scheme.

- Example 6a follows *Given-Given* ordering — The object “Amar Ujala” in the sentence is mentioned in the preceding context sentence 5, it would be annotated as *Given*. In contrast, the subject “yah” is a pronoun so it would also be tagged as *Given* following the annotation scheme.
- Example 6c follows *New-Given* ordering — The object “sukravar” in the sentence should be tagged as *New* as it is not mentioned in the preceding context sentence 5. In contrast, the subsequent pronoun “yah” which acts as the subject of the sentence should be tagged as *Given* following the annotation scheme.

C Adaptation Learning Rate

Table 5 illustrates the results of our learning rate experiments. Interestingly, van Schijndel and Linzen (2018) found that an adaptive learning rate of 2 minimized validation perplexity in English as well, though we leave further investigation of this to future work.

Learning Rate	0	0.002	0.02	0.2	2	20	200
Perplexity	103.29	98.79	87.78	66.64	56.86	117.91	$\sim 10^9$

Table 5: Learning rate influence on lexical and syntactic adaptation for the validation set containing 13274 sentences (the initial non-adaptive model performance is when we use a learning rate of 0)

D Correlation Plot

The Pearson’s correlation coefficients between different predictors are displayed in Figure 2. The adaptive LSTM surprisal has a high correlation with all other surprisal features and a low correlation with dependency length and information status score.

E GIVE Verb Class Regression Model

Predictor	$\hat{\beta}$	$\hat{\sigma}$	t
intercept	1.50	0.002	638.32
trigram surprisal	-0.11	0.013	-8.57
dependency length	0.01	0.003	2.78
pcfg surprisal	-0.08	0.004	-18.87
IS score	0.02	0.002	10.01
lex-rept surprisal	0.01	0.012	0.46
lstm surprisal	0.08	0.036	2.25
adaptive lstm surprisal	-0.36	0.037	-9.86

Table 6: Regression model on lemma verb GIVE data set (14094 data points; all significant predictors denoted by $|t|>2$)

⁸Higher and lower cutoffs do not affect our results.




Figure 2: Pearson’s coefficient of correlation between different pairs of predictors

F Levin’s Verb Class and Case Density

Verb Types	Case density	Freq	Freq (%)
GIVE	0.45	372	18.64
DO	0.39	726	36.37
COMMUNICATION	0.67	264	13.23
MOTION	0.39	93	4.66
SOCIAL	0.4	242	12.12
PERCEPTION	0.32	36	1.8
DESTROY	0.63	34	1.7
LODGE	0.32	95	4.76
PUT	0.4	52	2.61
OTHERS	0.43	82	4.11
Full	0.44	1996	100

Table 7: Levin’s verb semantic classes and case density (i.e., number of case markers per constituent in a sentence)

G Argument Ordering and Case Density

Alternation	Case density	Freq	Freq (%)
S-IO-DO	0.48	185	9.27
S-DO	0.39	1417	70.99
S-IO	0.59	394	19.74
Full	0.44	1996	100

Table 8: Argument ordering and case density (i.e., number of case markers per constituent in a sentence)

H Levin’s classes of verbs within Double Object (S-IO-DO) alternation

Verb Lemma	Frequency	Freq (%)	Verb Types	Freq (%)
<i>chah</i>	127	1.37	SOCIAL	2.59
<i>nawaja</i>	5	0.05		
<i>mil</i>	5	0.05		
<i>bech</i>	104	1.12		
<i>daal</i>	99	1.07	PUT	2.13
<i>jutaa</i>	75	0.81		
<i>pilaa</i>	23	0.25		
<i>dikha</i>	28	0.3	PERCEPTION	0.3
<i>badal</i>	99	1.07	LODGE	1.07
de	3240	34.92	GIVE	57.82
<i>saup</i>	1090	11.75		
<i>bhej</i>	569	6.13		
<i>maang</i>	419	4.52		
<i>dilaa</i>	46	0.5		
<i>kar</i>	1737	18.72		
<i>karaa</i>	465	5.01	DO	24.03
<i>chipaa</i>	23	0.25		
ban	5	0.05		
<i>kah</i>	883	9.52		
<i>sunaa</i>	198	2.13	COMMUNICATION	12.06
<i>likh</i>	23	0.25		
<i>bataa</i>	15	0.16		
Full (S-IO-DO)	9278	100		

Table 9: Levin’s syntactico-semantic classes of verbs within S-IO-DO data points from Table 1

I Information Profile: Syntactic Priming

Type	Trigram surp	Deplen	PCFG surp	IS score	LSTM surp	Adaptive LSTM surp	Lex rept surp
Example 2a Reference	34.27	24	107.04	0	173.06	156.88	36.45
Example 2b Variant	33.92	23	105.11	0	171.49	165.86	36.45
Example 4a Reference	58.04	40	144.98	-1	186.10	185.75	54.43
Example 4b Variant	57.68	26	143.06	1	185.31	184.52	56.84

Table 10: Predictor scores for reference-variant pairs




Figure 3: Information profiles for the reference-variant pair 2a and 2b

J Broad Coverage Analysis

Our regression results over the entire data set (Table 11) indicate that all the measures considered in our work are significant predictors of syntactic choice (*i.e.*, classifying reference and variant sentences). The negative regression coefficients for all surprisal metrics indicate that log-odds of predicting the reference sentences increase with decrease in their surprisal values. In other words, corpus reference sentences have consistently lower surprisal scores compared with the artificially generated competing variants. And adding adaptive LSTM surprisal into a model containing all other predictors significantly improved the fit of our regression model ($\chi^2 = 66.81$; $p < 0.001$). The positive regression coefficient for information status (IS) score indicates that reference sentences adhere to *given-new* ordering. Similarly, adding IS score into a model containing all other predictors significantly improved the fit of our regression model ($\chi^2 = 127.94$; $p < 0.001$). However, the positive regression coefficient of dependency length suggests that reference sentences exhibit *longer* dependency lengths compared to their variant counterparts, violating locality considerations. This further conjectures that dependency length might be in conflict with (and/or overridden by) other factors like discourse and priming. Future work needs to investigate if word-order preferences can be jointly optimized using multiple factors (Gildea and Jaeger, 2015).

We now examine the relative performance of each predictor in classifying reference sentences against the paired counterfactual grammatical variant by estimating the prediction accuracy (*i.e.*, the percentage of data points where the model chose the reference sentence as the best choice compared to the paired variant). We performed 10-fold cross-validation, trained the model on 9 folds, and generated its prediction on the remaining fold. Table 12 presents the individual as well as collective prediction performance of our predictors. Among individual predictor performances (Left side of Table 12; Full data), both adaptive and non-adapt LSTM surprisal achieved the highest classification accuracy. However, over a baseline

Predictor	$\hat{\beta}$	$\hat{\sigma}$	t
intercept	1.50	0.001	1496.47
trigram surprisal	-0.08	0.005	-14.53
dependency length	0.02	0.001	15.55
pcfg surprisal	-0.07	0.002	-39.46
IS score	0.01	0.001	11.32
lex-rept surprisal	-0.03	0.005	-5.31
lstm surprisal	-0.14	0.016	-9.26
adaptive lstm surprisal	-0.13	0.016	-8.18

Table 11: Regression model on full data set ($N = 72833$; all significant predictors denoted by $|t|>2$)

Predictors	Full Accuracy %	Conjunct Verb	Predictors	Full Accuracy %	Conjunct Verb
a = IS score	51.84	52.08	Collective: with repetition effects		
b = dep length	62.31***	66.32***	base1 = a+b+c+d+e+f	95.05	96.33
c = pcfg surp	86.86***	89.20***	base1 + g	95.06	96.34
d = lex repetition surp	90.07***	92.69***	Collective: without repetition effects		
e = 3-gram surp	91.18***	93.54***	base2 = a+b+c+e+f	95.06	96.34
f = lstm surp	94.01***	95.67***	base2 + g	95.09*	96.38*
g = adaptive lstm surp	94.06	95.68			

Table 12: Prediction performances (Full data set (72833 points), Conjunct Verb (51617 points); each row refers to a distinct model; *** McNemar's two-tailed significance compared to model on previous row)

model comprising every other predictor, adaptive LSTM surprisal induced a significant boost of 0.03% in classification accuracy ($p = 0.04$ using McNemar's two-tailed test) only when lexical repetition surprisal was not included in the model.