Degrees of animacy and telicity in Hindi intransitives: A corpus study

Eva Neu, Rajesh Bhatt & Shaunak Phadnis*

Abstract. Intransitives are known to fall into two classes, unergatives and unaccusatives, associated with a wide range of diverging properties. It has long been argued that the semantics of a given verb determines whether it will preferentially exhibit unergative or unaccusative behavior. In particular, agentive semantics has been associated with unergativity, telic semantics with unaccusativity. In this paper, we test this claim empirically using corpus data from Hindi. For a wide range of intransitives, we compute how often they take an animate argument and how often they combine with a telic light verb. The results offer some support for the effect of agentivity and telicity on the syntax of intransitives, but are too limited in scope to allow for reliable conclusions. We end by discussing the methodological challenges of establishing the semantic correlates of split intransitivity.

Keywords. unaccusativity; unergativity; telicity; agentivity; Hindi

1. Background. It has long been observed that intransitive verbs do not form a single homogeneous class but show a split in their behavior with respect to a wide range of diagnostics. According to the Unaccusative Hypothesis, originating with Burzio (1981, 1986) and Perlmutter (1978), intransitives can realize two distinct underlying structures. In an unergative syntax, the single argument is base-generated as an external argument of a functional head, labeled here v, where it is assigned an agent θ -role (1a).¹ In an unaccusative syntax, on the other hand, the subject argument originates as an internal argument in the complement position of the verb receiving a patient or theme θ -role (1b). In both cases, the argument subsequently moves to the subject position, SpecTP.



While the structures in (1a) and (1b) can thus be string-identical on the surface, a wide variety of constructions have been proposed to diagnose a difference between the two. For instance, reduced relative clauses are argued to only be licensed with an internal argument and thus to signal

^{*} Thanks to Songhee Kim, Kyle Johnson and Brian Dillon. Authors: Eva Neu, University of Massachusetts Amherst (eneu@umass.edu), Rajesh Bhatt, University of Massachusetts Amherst (bhatt@umass.edu) & Shaunak Phadnis, University of Massachusetts Amherst (sphadnis@umass.edu).

¹ The description is anachronistic in that the Unacusativity Hypothesis of course predates the now widely adopted Split VP Hypothesis (e.g., Kratzer 1996; Marantz 1997).

an unaccusative syntax (Embick 2004) and impersonal passives to require an external argument, only being compatible with an unergative syntax (Perlmutter 1978).

The claim that unergative and unaccusative verbs differ in their underlying structure, as schematized in (1), is not uncontroversial, and it has been argued that unaccusativity/unergativity diagnostics such as reduced relatives and impersonal passives can be accounted for solely in semantic terms (e.g., Van Valin 1990). However, the observation that intransitives cluster around two distinct types, however defined, is hardly disputed. For the sake of concreteness, we will assume a syntactic basis for unaccusativity in the following, but we believe that our study offers insights also under an analysis where the semantics of the verb directly determine its behavior with respect to unaccusativity/unergativity diagnostics, without any difference in the underlying syntax.

Split intransitivity is sometimes thought of as a split between two non-overlapping classes of verb types. For instance, English verbs such as *work*, *laugh* and *play* are commonly classified as unergatives, whereas verbs like (intransitive) *break*, *melt* and *fall* are considered typical unaccusatives. While often a convenient assumption in practice, this view has long been known to be an oversimplification, with many verbs allowing for both an unergative and an unaccusative use, passing unaccusativity diagnostics in some cases and unergativity diagnostics in others (e.g., Borer 2005; Perlmutter & Postal 1984). Sorace (2000, 2004, 2011) has argued that verbs lie on a spectrum from rigidly unaccusative to rigidly unergative, with verbs in the middle of the spectrum showing variable behavior.

Even under a syntactic analysis of unaccusativity, it is widely assumed that the semantics of the verb determine its position on this spectrum. In particular, two factors have been argued to push the verb towards one intransitive class or the other, namely, agentivity and telicity (e.g., Dowty 1991; Sorace 2000, 2004, 2011). To begin with the former, since the external argument position is associated with an agent θ -role, verbs that attribute proto-agent properties such as intentionality, purpose, volition, animacy/sentience and causal power to their argument tend to behave as unergatives. For instance, *work* describes a more agentive activity than *fall*, making it a better unergative. On the other hand, a typical property of patient arguments is to undergo a change of state. Accordingly, verbs describing a telic event, where the argument ends up in a specific resultant state, are more likely to behave as unaccusatives. For instance, the argument of the typical unaccusative *melt* undergoes an obvious change of state – from solid to liquid –, whereas the argument of the typical unergative *laugh* does not.

Agentivity and telicity have been argued to not only determine the preferred behavior of a given verb type but also to affect the behavior of different tokens of this type. For instance, it is well known that some verbs tend to behave as unergatives with animate arguments – animacy or sentience being one dimension of agentivity (Dowty 1991) – but as unaccusatives with inanimate arguments. In the Italian example in (2), this effect is demonstrated using auxiliary selection as a diagnostic, with unergatives selecting *have* and unaccusatives selecting *be*:

- (2) a. Il pilota ha/?è atterrato sulla pista di emergenza. the pilot has/is landed on.the runway of emergency 'The pilot landed on the emergency runway.'
 - b. L'aereo è/?ha atterrato sulla pista di emergenza. the.plane is/has landed on.the runway of emergency 'The plane landed on the emergency runway.'

(Sorace 2000:876)

The animate argument *pilot* in (2a), acting intentionally, is realized as an agent argument whereas the inanimate *plane* in (2b) is realized as a patient. As for telicity, (3) demonstrates that in Dutch, atelic *roll* shows unergative, but telic *roll* unaccusative behavior, again indicated by auxiliary choice:

- (3) a. De bal heeft/*is gerold. the ball has/is rolled 'The ball rolled.'
 - b. De bal is/*heeft naar beneden gerold. the ball has/is to down rolled 'The ball rolled downstairs.'

(Sorace 2000:876)

In (3b), the ball reaches the resultant state of being downstairs, thus qualifying as a patient rather than as an agent argument.

The claim that agentive and telic semantics affect a verb's syntactic behavior has enjoyed widespread acceptance for a long time, but the data it rests on have always been limited. For the most part, the effect has been demonstrated for isolated examples with a small number of verbs and based on the intuitions of a single researcher. Few reliable quantitative studies exist, and only on a limited sample of languages. Most recently, Kim et al. (2024) have carried out a large-scale experimental study, which we discuss below in Section 4; earlier work includes Acartürk & Zeyrek (2010); Allman (2017); Baker (2019); Huang (2018). A frequent methodological issue, pointed out by Kim et al. (2024), is that previous studies have often presupposed a) that judgments on unaccusativity diagnostics are dichotomous, and b) that agentivity and telicity are binary features, two assumptions that affect the way in which data are collected but that are by far not uncontroversial.

The present paper seeks to contribute to current efforts to remedy these limitations by presenting empirical measures of agentivity and telicity from Hindi corpus data. We assume that agentivity and telicity may be continuous rather than categorical semantic properties and also that syntactically, verbs may tend towards the unergative or the unaccusative end of the spectrum to varying degrees. Section 2 describes the methods of our study and Section 3 summarizes the results. In Section 4, we then discuss the methodological challenges of establishing the semantic correlates of split intransitivity.

2. Methods. Our measures of agentivity and telicity were developed for intransitive verbs in the Hindi-Urdu Treebank (Bhat et al. 2017). Hindi verbs are morphologically marked for valency, allowing us to specifically extract intransitives from the corpus, which makes Hindi a particularly suitable language for our purposes.

Since agentivity is a broad and loosely defined concept that is difficult to quantify, we focused on developing a measure for animacy, which is strongly correlated with agentivity more broadly and easier to operationalize. Specifically, we computed for each intransitive verb in the corpus how likely it is to appear with an animate or inanimate argument. To classify arguments based on animacy, we relied on case marking. In Hindi, arguments in subject position bearing ergative case are highly likely to be animate, whereas arguments in object position not bearing DOM marking are highly likely to be inanimate. On the basis of case marking, we thus compiled lists of animate and inanimate nouns, which we additionally checked manually. We then computed for each intransitive verb in the treebank how often its subject argument appears on the list of animate or inanimate nouns, respectively. We predict that verbs with a higher animacy score should be more likely to behave as unergatives.

To quantify telicity, we took advantage of the fact that Hindi has a variety of light verbs, two of which -a: 'come' or *ja*: 'go' - are associated with a telic interpretation. For each root, we computed how often it occurs with a telic light verb. To normalize these counts, we then divided them by the number of occurrences of the intransitive form in the corpus. We predict that verbs with a higher telicity score are more likely to behave as unaccusatives.

3. Results. Table 1 summarizes the results for intransitives with an overall corpus frequency of 5 or higher. The verbs are ordered based on their animacy score, from low to high. Note that some verbs, such as *bik* 'be sold' and *bigat* 'be spoiled,' can only be transitive in English but have intransitive variants in Hindi.

verb	Translation	Frequency	Animacy score	Telicity score (absolute)	Telicity score (normalized)
		_		2	
bik	be sold	5	0	0	0
sulajh	get untangled	5	0	2	0.4
chhaː	cover	6	0	0	0
bigaț	be spoiled	6	0	0	0
biːt	elapse	6	0	0	0
sudhar	improve	8	0	1	0.125
țuuț	break	15	0	12	0.8
khul	open	25	0.04	0	0
baṛh	increase	66	0.08	33	0.5
paŗ	fall	58	0.1	6	0.1
ghaț	reduce	9	0.11	0	0
phail	spread	9	0.11	0	0
ruk	stop	9	0.11	4	0.4
tik	stay	8	0.125	0	0
bah	flow	7	0.14	8	1.14
jam	freeze	5	0.2	6	1.2
bandh	be tied	5	0.2	0	0
nikal	emerge	33	0.21	13	0.39
gir	fall	9	0.22	10	1.11
ho	be	12221	0.23	755	0.06
chal	walk, function	134	0.28	9	0.07
jur	be connected	18	0.33	4	0.22
ban	become	136	0.36	70	0.51
bach	be saved	24	0.375	11	0.46
haːr	lose	5	0.4	4	0.8
ghir	be surrounded	5	0.4	0	0
baith	sit	15	0.47	0	0
ubhar	emerge	6	0.5	0	0
gujar	elapse	8	0.5	2	0.25

mar	die	8	0.5	8	1
pahunc	arrive	107	0.5	61	0.57
jaː	go	111	0.55	0	0
lauț	return	21	0.57	16	0.76
dauŗ	run	5	0.6	4	0.8
bha:g	run, escape	8	0.75	8	1
haț	remove oneself	6	1	1	0.17

Table 1. Summary of results

As for animacy, verbs towards the top of the table with the lowest animacy scores are predicted to be unaccusative-leaning, and those towards the bottom are predicted to be unergativeleaning. By and large, these predictions appear to be borne out. Verbs like *ja:* 'go' (0.55), *daur* 'run' (0.6) and *bha:g* 'run, escape' (0.75) are assumed to predominantly behave as unergatives, verbs like *tuut* 'break' (0), *khul* 'open' (0.004) and *par* 'fall' (0.1) as unaccusatives. Overall, animacy scores seem to be reasonably predictive of the syntactic behavior of intransitives.

The results for telicity are less clear-cut. On the one hand side, some typically unaccusative verbs achieve high telicity scores as predicted, e.g., *jam* 'freeze' (1.2), *tuut* 'break' (0.8) and *mar* 'die' (0.5). Equally, some typically unergative verbs score low, e.g., *jat* 'go' (0). However, some verbs typically considered unergatives receive high telicity scores, such as *laut* 'return' (0.76), *daur* 'run' (0.8) and *bhatg* 'run, escape' (1). For the latter two, it should be noted that verbs of directed motion are known to sometimes switch to an unaccusative use when used in a telic context, as shown above for Dutch in (3) and below for Hindi in (4). While atelic 'run' cannot occur in a reduced relative clause, which requires an unaccusative syntax, telic 'run' can appear in such environments for some speakers.

(4)	a.	# dauṛ-aː laṛkaː	b.	%sku:l=tak dauṛ-a: laṛka:
		run-PFV boy		school=to run-PFV boy
		*'the run boy'		'the boy who ran to school'

Thus, high telicity scores for verbs of directed motion might accurately reflect the fact that these verbs can receive a telic interpretation in contexts such as (4b), then adopting an unaccusative use.

Finally, some unaccusative-leaning change-of-state of verbs receive a low telicity score, among them *khul* 'open' (0), *ghat* 'reduce' (0) and *sudhar* 'improve' (0.125). It is not clear whether this is due to noise in the data – 'reduce' and 'improve' in particular have low absolute frequencies –, the potentially open-ended nature of the change of state described by these verbs or other factors. In sum, telicity scores do not always clearly line up with how verbs have been classified in the past. While this might partly reflect a greater flexibility in verbal behavior than often assumed, it might also point towards shortcomings of the empirical measure used here.

4. Discussion. To give a brief interim summary, the results for animacy have suggested a correlation with an unergative syntax, while those for telicity have been more complex but encouraging. However, the findings are subject to three major caveats. First, animacy is only an imperfect proxy for the broader concept of agentivity, which also includes dimensions like purpose and causal power. Second, corpus frequencies for the verbs in the sample are low overall; hence, the

results must be taken with a sizable grain of salt. In particular, we found only few occurrences of roots with telic light verbs, making the telicity measure especially prone to being distorted by noise. Third, in order to conclusively establish a correlation between semantic agentivity/telicity and syntactic unergativity/unaccusativity, we would need a quantitative measure of the latter as well. So far, we have only compared the corpus results to the unergative/unaccusative status of verbs as it has been reported in the literature, which might underestimate the degree to which verbs can vary between the two classes. Providing an empirical measure of the syntactic behavior of different intransitives would be desirable, but given that the semantic measures we computed are not based on sufficiently robust numbers, pursuing this line for Hindi does not appear to hold much promise.

In the remainder of this paper, we discuss different methodological avenues for testing the link between agentivity/telicity and unergativity/unaccusativity. The question we focus on is how to operationalize agentivity and telicity, which are after all latent semantic dimensions that verbs do not bear on their sleeves. We discuss two different approaches to this challenge, which we label introspective and distributional. As an example of the former approach, we briefly review recent work by Kim et al. (2024); the Hindi corpus study presented here is an example of the latter approach.

By introspective we refer to the strategy of having speakers directly access their intuitions about the degree of agentivity or telicity of a given verb. This strategy is the default in much work on the semantic correlates of unaccusativity/unergativity, which typically relies on the intuitions of a single researcher. To make such introspective measures more robust, Kim et al. (2024) have collected large-scale experimental data by asking participants to rate verbs for agentivity and telicity (besides other features) on a scale from 0 to 6.

While providing valuable evidence, there are also potential pitfalls to the introspective approach. First, in an experimental setting, much depends on the precise wording of the prompt. For instance, to test for agentivity, Kim et al. asked their participants to what extent a given verb describes 'something that is actively or intentionally done.' This question does not take into account animacy or sentience, often considered a core component of agentivity (Dowty 1991). The omission is reflected in Kim et al.'s results, with verbs like 'snore' and 'stumble,' which almost always take an animate argument but describe unintentional actions, receiving very low average scores below 1. A different wording of the question might have led to these verbs being rated very differently, for better or for worse. The point here is not that Kim et al.'s prompt is wrong, but simply that there is a non-trivial choice to make which will affect the findings.

The second challenge for introspective measures is that they are typically collected for verbs in isolation, as also done by Kim et al., or only in the context of a single sentence. However, agentivity and telicity are heavily influenced by their environment. As discussed earlier, verbs of directed motion like *run* can receive a telic interpretation in the context of modifiers like *to school* but an atelic interpretation otherwise. By the same token, a verb like *fly* denotes an intentional and effortful activity with animate arguments such as *the bird*, but can also describe a state more akin to passive floating when it takes an inanimate argument like *the kite*.

In short, how agentive and how telic the event denoted by the verb is depends on the sentence this verb appears in. The question this raises is what intuition speakers are actually accessing when asked to quantify the degree of agentivity and telicity of a verb in isolation. It might be an average of the different usages that a verb permits, or – which is perhaps more likely – it might be the degree of agentivity and telicity of the event described by a prototypical use of the verb. For instance, when queried about the verb *run*, speakers might access their representation of a running event in a sentence like (5a), but not in sentences (5b)–(5e).

- (5) a. He is running.
 - b. He ran to the supermarket in just five minutes.
 - c. A Broadway show usually runs for several months.
 - d. Please keep the car running.
 - e. I need to run some code.

In other words, given that agentivity and telicity are always context-dependent, it is not obvious that intuitions about verbs in isolation give us a meaningful representation of these properties, nor that they will be predictive of a verb's syntactic behavior.

A very different approach to quantifying agentivity and telicity is to rely on distributional methods. Here the strategy is to quantify certain properties of the contexts in which individual tokens of the verb occur. For instance, in this paper we have counted the frequencies with which a verb occurs in the context of an animate argument or a telic light verb. This approach does not presuppose that the verb has a single, unified semantic core that can be described as having a certain degree of agentivity and telicity. Rather, the verb is regarded as the total sum of its usages, and the distributional measures of agentivity and telicity constitute averages over those usages. This avoids the second problem of the introspective approach outlined earlier, which is to judge a verb in isolation from its context. It does not, of course, solve the first problem. As before, there is no single unambiguous contextual correlate of agentivity or telicity. For instance, the animacy score we computed is clearly not an exhaustive measure of agentivity, and different measures would yield different results.

There is an additional potential problem for at least some distributional measures. Even if we could conclusively establish a correlation between the likelihood of a verb to take animate arguments and to appear in an unergative structure – presupposing that the latter can be quantified in some way –, this might also be taken to simply indicate a direct correlation between animate arguments and an unergative syntax, without a root middleman. In this way, we might establish a correlation between animacy and unergativity, but it is not clear how exactly the semantics of the root would enter into the equation.

The broader question this raises is whether agentivity and telicity should be understood as semantic properties of verbs in the first place. Telicity in particular has long been argued to rather be a property of verb phrases (Borer 2005; Rothstein 2008; Tenny 1987, a.m.o.), and agentivity can equally be thought of as describing an entire event description, being modulated by the verb just as much as by its arguments. Of course, different verbs affect these VP- or sentence-level properties in different ways; all other things being equal, *work* does describe more agentive events than *fall*. But testing whether a sentence with a strongly agentive and weakly telic semantics is more likely to have an unergative syntax, and vice versa for an unaccusative syntax, might be a more fruitful avenue for future research than attempting to pin down the contribution of the verb itself.

5. Conclusion. This paper has tested the long-standing claim that agentive semantics is correlated with an unergative syntax, and telic semantics with an unaccusative syntax. Using corpus data from Hindi, we have investigated how often different intransitives take an animate argument and combine with a telic light verb. We have found moderate support for the hypothesis, but ulti-

mately our data set has proven too small to allow for reliable conclusions.

Against this background, we have discussed and compared two different methodological approaches to testing the semantic correlates of split intransitivity, namely, collecting intuitions about the verb itself and quantifying properties of its context. We hope to have shown that agentivity and telicity are not straightforward to operationalize. Both concepts – agentivity in particular – are notoriously broad and multi-faceted, and much depends on the wording of the experimental prompt or the concrete contextual measure. What is more, it is not obvious that agentivity and telicity can be thought of as properties that a verb possesses in isolation from its context. We have suggested that ultimately, investigating the correlation between agentivity/telicity and the unergative/unaccusative split at the level of the sentence might hold more promise.

References

- Acartürk, Cengiz & Deniz Zeyrek. 2010. Unaccusative/unergative distinction in Turkish: A connectionist approach. In *Proceedings of the 8th workshop on Asian language resouces*, 111–119. Beijing, China: Coling 2010 Organizing Committee.
- Allman, JungAe Lee. 2017. *Empirical examination of two diagnostics of Korean unaccusativity*. Arlington, TX: The University of Texas at Arlington dissertation.
- Baker, James. 2019. Split intransitivity in english. *English Language and Linguistics* 23. 557589. https://doi.org/10.1017/S1360674317000533.
- Bhat, Riyaz Ahmad, Rajesh Bhatt, Annahita Farudi, Prescott Klassen, Bhuvana Narasimhan, Martha Palmer, Owen Rambow, Dipti Misra Sharma, Ashwini Vaidya, Sri Ramagurumurthy Vishnu et al. 2017. The Hindi/Urdu Treebank project. In Pustejovsky J. Ide, N. (ed.), *Handbook of linguistic annotation*, Springer Press. https://doi.org/10.1007/978-94-024-0881-2₂4.
- Borer, Hagit. 2005. *Structuring sense vol. II: The normal course of events*. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199263929.001.0001.
- Burzio, Luigi. 1981. Intranitive verbs and Italian auxiliaries. Cambridge, MA: dissertation.
- Burzio, Luigi. 1986. *Italian syntax: A Government and Binding approach*. Dordrecht: D. Reidel. https://doi.org/https://doi.org/10.1007/978-94-009-4522-7.
- Dowty, David. 1991. Thematic proto-roles and argument selection. *Language* 67. 547–619. https://doi.org/10.2307/415037.
- Embick, David. 2004. Unaccusative syntax and verbal alternations. In Artemis Alexiadou, Elena Anagnostopoulou & Martin Everaert (eds.), *The unaccusativity puzzle: Explorations of the syntax-lexicon interface*, 137–158. Oxford: Oxford UP. https://doi.org/10.1093/acprof:oso/9780199257652.003.0006.
- Huang, Yujing. 2018. *Linking form to meaning: Reevaluating the evidence for the unaccusative hypothesis.* Cambridge, MA.
- Kim, Songhee, Jeffrey R Binder, Colin Humphries & Lisa L Conant. 2024. Decomposing unaccusativity: a statistical modelling approach. *Language, Cognition and Neuroscience* 39. 1189–1211. https://doi.org/10.1080/23273798.2024.2368119.
- Kratzer, Angelika. 1996. Severing the external argument from its verb. In Johan Rooryck & Laurie Zaring (eds.), *Phrase structure and the lexicon*, 109–137. Kluwer/Springer. https://doi.org/10.1007/978-94-015-8617-7₅.
- Marantz, Alec. 1997. No escape from syntax: Dont try morphological analysis in the privacy of your own lexicon. In *University of Pennsylvania working papers in linguistics* 4.2, 201–225.

- Perlmutter, David. 1978. Impersonal passives and the Unaccusative Hypothesis. *Papers from the Annual Meeting of the Berkeley Linguistic Society* 4. 157–189. https://doi.org/10.3765/bls.v4i0.2198.
- Perlmutter, David & Paul Postal. 1984. The 1-Advancement Exclusiveness Law. In David Perlmutter & Carol Rosen (eds.), *Studies in Relational Grammar*, 81–125. Chicago, IL: University of Chicago Press.
- Rothstein, Susan. 2008. *Structuring events: A study in the semantics of lexical aspect.* https://doi.org/10.1002/9780470759127.
- Sorace, Antonella. 2000. Gradients in auxiliary selection with intransitive verbs. *Language* 76(4). 859–890. https://doi.org/10.2307/417202.
- Sorace, Antonella. 2004. Gradience at the lexicon-syntax interface: Evidence from auxiliary selection and implications for unaccusativity. In Artemis Alexiadou, Elena Anagnostopoulou & Martin Everaert (eds.), *The unaccusativity puzzle*, 243–268. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199257652.003.0010.
- Sorace, Antonella. 2011. Gradience in split intransitivity: the end of the Unaccusative Hypothesis? *Archivio Glottologico Italiano* XCVI(1). 67–86.

Tenny, Carol. 1987. Grammaticalizing aspect and affectedness. Cambridge, MA: dissertation.

Van Valin, Robert D. 1990. Semantic parameters of split intransitivity. *Language* 221–260. https://doi.org/10.2307/414886.