

A Dataset for Physical and Abstract Plausibility and Sources of Human Disagreement

Annerose Eichel, Sabine Schulte im Walde

Institute for Natural Language Processing, University of Stuttgart
{annerose.eichel,schulte}@ims.uni-stuttgart.de

Abstract

We present a novel dataset for physical and abstract plausibility of events in English. Based on naturally occurring sentences extracted from Wikipedia, we infiltrate degrees of abstractness, and automatically generate perturbed pseudo-implausible events. We annotate a filtered and balanced subset for plausibility using crowd-sourcing, and perform extensive cleansing to ensure annotation quality. In-depth quantitative analyses indicate that annotators favor plausibility over implausibility and disagree more on implausible events. Furthermore, our plausibility dataset is the first to capture abstractness in events to the same extent as concreteness, and we find that event abstractness has an impact on plausibility ratings: more concrete event participants trigger a perception of implausibility.

1 Introduction

The ability to discern plausible from implausible events is a crucial building block for natural language processing (NLP). Most previous work on modelling plausibility however focuses on the kinds of semantic knowledge necessary for distinguishing a *physically* plausible event from an implausible one (Wang et al., 2018; Porada et al., 2019). As illustrated in Fig. 1, the current study extends the traditional focus to discern physically plausible events such as *cat-eat-sardine* from physically implausible ones such as *rain-break-belly*. Furthermore, while recent datasets include some events with conceptually *abstract* participants (Emami et al., 2021; Pyatkin et al., 2021), as to our knowledge no dataset nor model up to date has paid attention to the interaction of event plausibility and abstractness of the involved concepts. We propose to systematically examine plausibility across levels of abstractness, and distinguish between abstractly plausible events such as *law-prohibit-discrimination* and abstractly implausible ones such as *humour-require-merger*. We hypothesize that (i) plausible vs. implausible events can

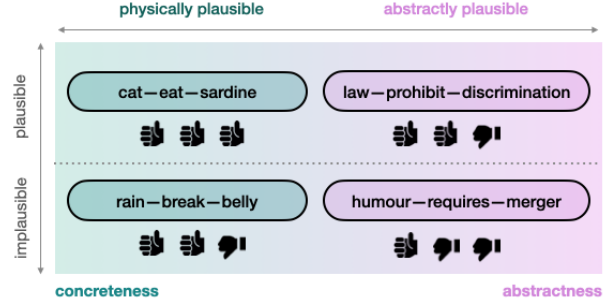


Figure 1: Plausible and implausible example events integrating degrees of concreteness/abstractness, cf. physical (green) and abstract (pink) levels. Annotators might agree (thumbs up) or disagree (thumbs down) on the (im)plausibility of the events.

be captured through physical vs. abstract levels, and that (ii) integrating degrees of abstractness into events fosters the understanding and modelling of plausibility (cf. Fig. 1).

We start out with a set of attested, i.e., *plausible*, natural language events in form of *s-v-o* triples from the English Wikipedia, assign abstractness ratings to event participants, and partition the triples into bins with varying degrees of abstractness. We then automatically generate pseudo-*implausible* event triples and assign degrees of abstractness in a similar way. To obtain human plausibility ratings for each event triple, we conduct a crowd-sourcing annotation study. We collect and validate a total of 15,571 judgements amounting to an average of 8.9 ratings for 1,733 event triples.

Human intuition regarding the assessment of plausibility is, however, incredibly multi-faceted, highly individual, and not easily reproducible automatically (Resnik, 1993). In particular, boundaries between categories to be annotated or predicted might not necessarily be strictly true or false, i.e., either plausible or implausible, thus reflecting the true underlying distribution of non-deterministic human judgements with inherent disagreement about labels (Baan et al., 2022). Over

the past decade, a growing body of work has emphasized the need to incorporate such disagreement in NLP datasets to reflect a more realistic and holistic picture across NLP tasks (Plank et al., 2014; Aroyo and Welty, 2015; Jamison and Gurevych, 2015; Basile et al., 2021b; Uma et al., 2021a). Accordingly, we argue for the necessity to preserve and examine disagreement when annotating and modelling plausibility, and represent inherent disagreement in annotation in order to devise a range of silver standards for analysis and modelling. More specifically, we disentangle subjectivity from annotation error, limitations of the annotation scheme, and interface (Pradhan et al., 2012; Poesio et al., 2019), and examine disagreements in physical and abstract plausibility annotation.

Overall, we find that our annotators tend to favor plausibility over implausibility, and we observe stronger disagreements for implausible in comparison to plausible events. Second, we explore the impact of abstractness on plausibility ratings. Here, our results reveal a positive relation between plausibility and events consisting of more abstract words, while implausibility is mostly found in predominantly concrete events.

2 Background and Related Work

2.1 Capturing (Semantic) Plausibility

The notion of plausibility has been approached from many perspectives. Inspired by the overview in Porada et al. (2021), we present distinctions and discuss viewpoints from previous work. Similarly to related notions such as *selectional preference* (Wilks, 1975; Resnik, 1993; Erk et al., 2010; Van de Cruys, 2014; Zhang et al., 2019; Metheniti et al., 2020) and *thematic fit* (Chersoni et al., 2016; Sayeed et al., 2016; Pedinotti et al., 2021), plausibility estimations capture non-surprisal in a given context. For example, the event *kid-sleep* with the agent *kid* is less surprising than *tree-sleep* and therefore considered more plausible. Within the context of (semantic) plausibility, however, plausible events are not necessarily assumed to be the most typical or preferred events. This stands in contrast with selectional preference or thematic fit, where whatever is not preferred is considered atypical albeit, in principle, a given event might be plausible. Wilks (1975) also discusses naturally occurring cases where the most preferred option does not yield the only correct interpretation: “[t]he point is to prefer the normal, but to *accept* the unusual.”

In this vein, Wang et al. (2018) propose the task of semantic plausibility as “recognizing plausible but possibly novel events”, where a “novel” event might be an unusual but nevertheless plausible event. Porada et al. (2021) further point out that “[p]lausibility is dictated by likelihood of occurrence on the world rather than text”, and attribute this discrepancy to reporting bias (Gordon and Van Durme, 2013; Schwartz and Choi, 2020). For example, it is much more likely that the event *human-dying* is attested than the event of *human-breathing*. The sum of all plausible events in a given world thus encompasses not only the sum of all attested events in a corpus (including modalities other than text), but also possibly plausible events which are *not* necessarily attested in a corpus.

In our definition what is *preferred* is considered the *most plausible*, but what is *unusual* might still be highly *plausible*. Plausibility therefore (i) **exceeds the boundaries of (selectional) preference**. Further, plausibility (ii) is **a matter of degree** as the preferred is considered *more* plausible. In turn, what is unusual is still considered plausible albeit to a lesser degree. Moreover, plausibility (iii) captures **non-surprisal in a given context**, and (iv) denotes what is generally **likely, but not necessarily attested in a given corpus**.

2.2 Measuring Semantic Plausibility

There are various positions on how to model, measure, and evaluate whether an event triple is plausible. In this study, we model plausibility as the proportion of what is considered plausible, requiring a minimal label set of {implausible, plausible} (Wang et al., 2018). Note that a value regarding what is “true” is not involved in measuring plausibility. Consider the examples *eat-strawberry*, *eat-pebble*, and *eat-skyscraper*. Given our label set, the first two events would be considered plausible (even though they strongly vary in their degree of plausibility and likelihood to be attested in text with eating a strawberry considered more plausible than the less, but still plausible process of eating a pebble)¹, while the last event is physically implausible. Derived label sets such as {implausible, neutral, plausible} may include a “neutral” label which is considered to not carry plausibility information, as it does not provide insight into whether an expression is (im)plausible (Anthonio et al., 2022).

¹Using e.g., Google n-grams, *eat-strawberry* is clearly attested more often than *eat-pebble*, while *eat-skyscraper* is not attested at all.

When annotating plausibility, drawing hard lines between labels is difficult and increases in complexity when considering words and concepts that are more abstract than concrete. This is especially true when considering free-standing events where no information on limiting factors regarding interpretation can be inferred. An example would be *human-breathe* which is plausible unless the human in question is dead. A more complex example would be *human-have-human_rights*, which is likely to be considered plausible by the majority of people and mirrored by corresponding laws in many countries, but (a) not universally accepted by each individual, and (b) not formalized as such by all countries.

2.3 Physical and Abstract Plausibility

Concepts can be described in accordance with the way people perceive them. While concepts that can be seen, heard, touched, smelled, or tasted are described as *concrete*, those that cannot be perceived with the five senses are referred to as *abstract* (Barsalou and Wiemer-Hastings, 2005; Brysbaert et al., 2014). Examples of concrete concepts include *apple*, *house* and *trampoline*, abstract examples encompass *absurdity*, *luck*, and *realism*. While instances at each extreme of abstractness occur, the notion is not binary but rather continuous, including many concepts between each extreme. Mid-range examples include concepts such as *inflation*, *punctuality* and *espionage*.

The grounding theory of cognition argues that humans process abstract concepts by creating a perceptual representation that is inherently concrete as it is generated through exposure to real world situations using our five senses (Van Dam et al., 2010; Brysbaert et al., 2014). However, more recent work brings forth evidence suggesting that such representations incorporate both perceptual and non-perceptual features (Dove, 2009; Naumann et al., 2018; Frassinelli and Schulte im Walde, 2019).

Regarding suitable abstractness ratings, we find a variety of datasets of growing size and diversity for many languages.² A widely used collection are the concreteness norms devised by Brysbaert et al. (2014), who collected ratings for approx. 40K “generally known” English words such as *sled* and *dream*, referring to strength of sense perception.

²See for a detailed overview, e.g., Schulte im Walde and Frassinelli (2022); Charbonnier and Wartena (2019).

2.4 Disagreement in Dataset Construction

While humans excel at assessing plausibility, they might naturally disagree regarding the plausibility of an event such as *law-prohibit-discrimination*. In the course of the last decade, a growing line of research argues for the preservation and integration of disagreement in dataset construction, modelling, and evaluation (Aroyo and Welty, 2015; Pavlick and Kwiatkowski, 2019; Basile et al., 2021b; Fornaciari et al., 2021; Uma et al., 2021a)³. While highly subjective tasks such as sentiment analysis (Yin et al., 2012; Kenyon-Dean et al., 2018) and offensive language detection (Leonardelli et al., 2021; Almanea and Poesio, 2022) have gathered particular attention, prior work has also presented evidence for seemingly objective tasks requiring linguistic knowledge such as PoS tagging (Gimpel et al., 2011; Hovy et al., 2014; Plank et al., 2014). We thus argue for the necessity to disentangle, devise, and examine disagreement when annotating and modelling plausibility. In contrast to previous work on plausibility assessments, we represent inherent disagreement in annotation and devise a range of silver standards for analysis and modelling.

3 Construction of Event Targets

Our first goal is to create a dataset⁴ that systematically (a) covers both plausible event triples that are selectionally preferred or unusual, (b) captures events attested in the real world, i.e., extracted from triples produced in natural language, (c) measures plausibility on a degree scale from plausible to implausible, and (d) puts equal emphasis on both abstractly and physically plausible events. We visualize the dataset construction process in Fig. 2.

3.1 Extracting Natural Language Triples

To compile a set of natural language triples, we first extract all text from an English Wikipedia dump using gensim (Řehůřek and Sojka, 2010). We then randomly sample k articles⁵ with $k=50,000$ and syntactically parse the articles using stanza (Qi et al., 2020). Next, we extract a triple (s, v, o) whenever the following conditions are satisfied: s

³For an overview we refer to, e.g., Basile et al. (2021a); Uma et al. (2021b).

⁴The dataset including event triples, ratings, and aggregated labels is available at <https://github.com/AnneroseEichel/PAP>.

⁵We only store the section texts and discard (section) titles as they tend to not contain 3-tuples of the desired form.

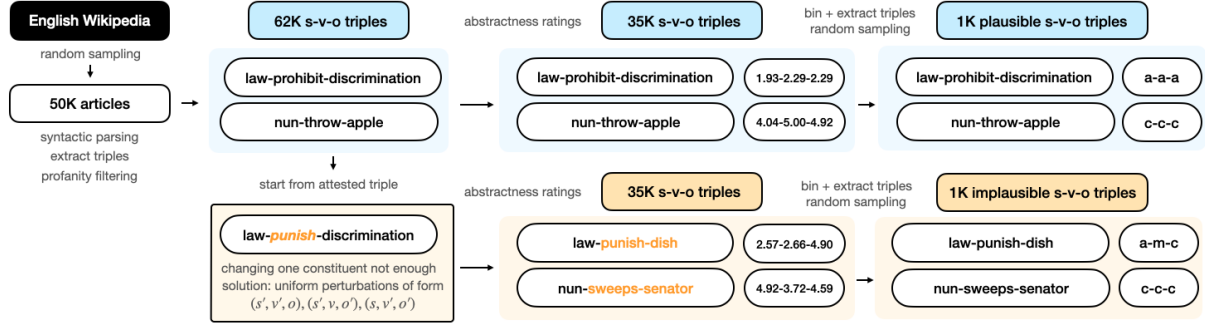


Figure 2: Simplified illustration of dataset construction starting with the extraction of attested event triples from a sample of the English Wikipedia. We filter triples, assign abstractness ratings, bin, and sample 1,080 plausible event triples for 27 abstractness combinations (marked in blue). Based on attested triples, we automatically generate pseudo-implausible triples and similarly filter triples, assign abstractness ratings, perform binning, and sample 1,080 implausible event triples (marked in yellow).

is the lemma of the head of *nsubj*, *o* is the lemma of the head of *obj*, and *v* is the lemma of the head of the root verb. We only allow nouns in subject and object positions and disregard proper names and pronouns as well as nouns and verbs that are part of a compound, yielding 62,843 triples. We extract each triple once, keeping track of frequency w.r.t sampled text data. Triples containing nouns or verbs that are explicit or have offensive connotations are filtered out using existing tools.⁶ In total, this leaves us with 62,473 triples.

3.2 Creating Physically and Abstractly Plausible Triples

To discern triples containing highly concrete words from triples which encompass more abstract words, we assign abstractness scores to all nouns and verbs in a triple, drawing on the concreteness ratings by Brysbaert et al. (2014). We use a reduced collection⁷ encompassing 12,880 noun and 2,522 verb targets to assign concreteness ratings to all 62,473 triples where a rating *r* exists for each word $w \in \{s, v, o\}$. Instances with nouns or verbs for which no rating exists are discarded. Overall, the assignment step yields 35,602 triples⁸ with ratings. As we are specifically interested in distinctive features of abstractness vs. concreteness and cases which can be found in the middle of the continuous scale, we partition each constituent and each triple into 5 bins [*highly abstract*, *abstract*, *mid-range*, *concrete*, *highly concrete*]. To construct our dataset, we then only consider the bins at each extreme as

well as the mid-range bin. Each constituent of a triple *t* can be either *highly abstract* (a), *mid-range* (m), or *highly concrete* (c). Taking the Cartesian product, we thus define 27 possible triple combinations, e.g., triples consisting of words with very high concrete ratings only, e.g., (c, c, c) or fully mixed triples, e.g., (c, m, a). To extract triples satisfying the conditions of each of the 27 possible triple combination, we carry out the following steps:

- Partition each constituent in *s, v, o* in each $triple_{1...n}$ into 5 bins of equal size, ranging from very abstract to very concrete. Whenever the relative threshold θ between bins prohibits perfectly equal sizes, we trade perfect bin size for perfectly separated abstractness ratings.
- Extract all triples satisfying the conditions of a combination e.g., (c, c, c) from our set of 35,602 triples.

The distribution of all naturally occurring triples for each triple combination $\in \{(a, a, a), \dots (c, c, c)\}$ is presented in Fig. 6, App. A.2. To select plausible triples for annotation, we randomly sample 40 triples for each combination, yielding a total of 1,080 plausible triples.

3.3 Constructing Physically and Abstractly Implausible Triples

To construct implausible triples, we use the 35,602 cleaned triples for which an abstractness rating as provided by Brysbaert et al. (2014) exists. This restriction makes the task of implausible triple generation non-trivial as the set of possible constituents in each function is now limited to subjects,

⁶filter-profanity, alt-profanity-check

⁷For details on the filtering process, cf. App. A.1.

⁸Subject and object types amount to 4,140 and 4,551 unique words, respectively, while verb types are significantly less diverse (1,218 unique words).

verbs and objects that are attested to be plausible in their given function. Generating perturbations of attested triples as used by Porada et al. (2021) –where only one constituent, e.g., the subject, is perturbed while verb and object are kept– also results in disproportionately many plausible triples, e.g., *jurisdiction-evaluate-reaction*.

We thus use only the following perturbations: For each $t \in$ attested triples, we obtain a randomly perturbed t' serving as a pseudo-implausible natural language triple. We uniformly generate perturbations of the form (s', v', o) , (s', v, o') and (s, v', o') , where s' , v' , and o' are arguments randomly sampled from the plausible triple collection taking into account corresponding functions, e.g., only words for which the use as object is attested in the corpus are randomly sampled as an object perturbation. We discard all triples that exist in the plausible triple collection and only keep unique instances, thus yielding 35,600 pseudo-implausible triples. After profanity filtering, we are left with 35,447 triples. We assign abstractness ratings and apply the binning method as described in the previous section 3.2.

The distribution of physically and abstractly pseudo-implausible triples per combination is shown in Fig. 6 (b), App. A.2. In analogy to plausible triple construction, we sample 40 triples for each abstractness combination to obtain 1,080 implausible triples.

4 Human Annotation

Our second goal targets the annotation of the collected event triples with respect to subjective assessments of plausibility on a degree scale (1–5) ranging from implausible to plausible. For this, we perform a human annotation study.

4.1 Collecting Ratings for (Im)Plausibility

Task We collect plausibility judgements on Amazon Mechanical Turk⁹ for our 2,160 plausible and implausible triples. Each triple is annotated by 10 annotators. In particular, we ask annotators to indicate whether a given sentence is implausible or plausible using a sliding bar (corresponding to a scale from 1 to 5). An example of the task with full instructions as presented to annotators in our Human Intelligence Task (HIT) is illustrated in Fig 7, App. B.1. To avoid bias, the slider is by default set to the middle of the bar. Annotators are required

to move the slider and thereby make a decision for either plausible or implausible. Task instructions clearly inform about the possibility of submission rejections if the slider remains in the middle position.

Annotators Participation is limited to annotators based in the United States and the United Kingdom. We further require annotators to have a HIT Approval Rate $> 98\%$ and a number of $\geq 1,000$ approved HITs from previous work.

Quality Checks To track annotation quality, we use an initial set of 20 manually produced check instances (cf. App. B.2) that were judged clearly plausible/implausible by the authors and an additional English native speaker. Annotators are presented batches of 24 randomly shuffled plausible or implausible triples, plus one randomly sampled check instance. In case of failed check instances, we discard all annotations submitted by the corresponding worker.

4.2 Annotation Post-Processing

After discarding submissions where the slider is set to the default (rating=3) as well as submission from workers who failed a check instance, we collect a total of 21,317 plausibility ratings.¹⁰ We further perform the following post-processing steps in order to minimise the impact of spam and low-quality annotations regarding the plausibility of a given event (Roller et al., 2013; Rodrigues et al., 2017; Leonardelli et al., 2021), with datasets statistics at every processing step shown in Table 3, App. B.3. We first filter out ratings from workers who submitted annotations for < 10 instances. Assuming that events observed in Wikipedia represent plausible events, we then exclude ratings from workers whose annotations disagree with the original label *plausible* in more than 75% of their corresponding submissions.

After these steps, our number n of annotators A still amounts to a large set of $nA > 500$ annotators. To ensure sufficient agreement between annotators, we calculate a *soft* pairwise Jaccard Coefficient J (Jaccard, 1902)¹¹ for all annotator combinations, and only keep annotations from workers whose submissions yield an average $J > 0.4$, following

¹⁰After our main collection round, we experiment with basic post-processing steps and perform a small second collection round to collect 5 ratings for 84 triples with < 8 ratings.

¹¹See App. B.3 for details on the calculation.

⁹<https://www.mturk.com/>




Figure 3: (a) Number of plausibility ratings per rating option where ratings below 3 denote implausibility and ratings above 3 denote plausibility. (b) Number of triples across ratings aggregated as averaged median ratings. Ratings range from implausible $\{1, 2\}$ to plausible $\{4, 5\}$.

Bettinger et al. (2020). Finally, we keep only triples in the dataset if they received at least 8 ratings.

4.3 Dataset Statistics

After post-processing, we are left with 15,571 plausibility ratings for 1,733 triples (80% of the original triple set). With respect to instance coverage per abstractness combination, we have an average number of 32 triples per combination for both plausible and implausible triples with a minimum of 27 triples for the combinations (a, a, m) and (m, c, c) for plausible and implausible triples, respectively. Triples receive between 8 and 12 ratings, with an average of 8.9 ratings.

Estimated average Inter-Annotator Agreement (IAA) across our post-processed dataset using the previously introduced soft pairwise Jaccard Coefficient reaches 0.64. This indicates reasonable agreement among annotators; cases of disagreement we will explore in the next section.

5 Analysis of Human Judgements and Disagreement

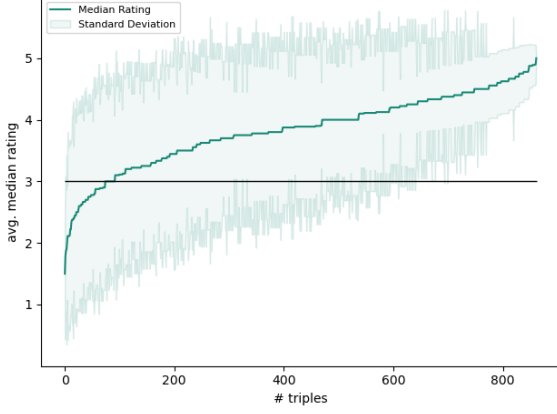
5.1 Examining Rating Distributions

Fig. 3 (a) shows the distribution of ratings across the four rating options, with green and pink bars indicating originally plausible and implausible label, respectively. The distribution is skewed towards plausibility with 68.98% ratings $\in \{4, 5\}$. We aggregate all individual ratings as average median rating per triple and show the resulting distribution in Fig. 3 (b). While the distribution for originally plausible triples (green bars) evens out as expected


with a peak number of average median rating for average plausibility (avg. median ratings $\in (3; 4]$), a similar peak can be observed given the distribution for originally implausibly triples (pink bars). The graph also shows differences, namely substantially more triples with a median rating indicating weak implausibility (avg. median rating $\in (2; 3]$) for originally implausible triples. On the other hand, high plausibility (rating $\in (4; 5]$) is annotated for mostly originally plausible triples.

To further investigate the skew towards plausibility, we visualize the average median rating for originally plausible and implausible triples in Fig. 4. The plot also illustrates the standard deviation of the values as a cloud. We observe that annotator ratings tend to show more overlap for plausible triples, with standard deviation decreasing with higher plausibility. In contrast, rating triples labeled as implausible result in greater deviation from the average mean rating decreasing only with implausibility. Taking into account the black horizontal line at a median rating of 3, we clearly see that median ratings for originally plausible triples are mostly above the cut line, thus indicating an overlap with the original label. On the other hand, median ratings for originally implausible triples are mostly below the cut line, thus indicating a clash with the original label.

These observations suggest (i) that **humans favor plausibility over implausibility**, while avoiding the extreme on the plausibility end of the scale, and (ii) that **implausibility yields higher disagreement**, as annotators disagree more when rating triples that were originally labeled as implausible.



(a) Average median rating across plausible triples.



(b) Average median rating across implausible triples.

Figure 4: Average median ratings across originally plausible (a) and implausible (b) triples with standard deviation visualized as cloud around average rating lines. Triples are represented numerically on the x-axis. The black horizontal line denotes a median rating of 3. Average median ratings for *plausible* triples *below* the line disagree with the original label, while the opposite is true for average median ratings for *implausible* triples. Here, ratings *above* the line disagree with the original label.

5.2 Exploring the Impact of Abstractness on Plausibility Ratings


Abstractness at Event Level To assess the relation between degrees of abstractness for combinations of words and plausibility on physical and abstract levels, we first examine the proportion of plausibility ratings across triples from each of our 27 abstractness combinations. For this, we calculate a *strict* majority ($\geq 70\%$) for each triple. Whenever ratings do not point to a majority, i.e., 50% *plausible* vs. 50% *implausible*, we mark the triple as *unsure*. We present a visualization in Fig. 5 where green bars denote a strict majority of plausible ratings $\in \{4, 5\}$, pink bars refer to a strict majority of implausible ratings $\in \{1, 2\}$, and orange bars illustrate the lack of clear majorities.

For attested plausible triples, original label and proportional majority rating overlap in all cases. In only three cases we observe majority ratings proportions below 50%, namely for the mostly concrete combinations (c, c, m) , (a, c, c) , and (a, c, m) . In contrast, majority rating proportions are generally higher for more abstract combinations, e.g., (a, a, a) , (m, a, a) . While a very low average of majority ratings for implausibility (1.3) can be observed, an average of 26.2 is obtained for triples with no majority. These observations suggest that (i) implausibility is most likely assigned to triples with concrete words, inducing higher disagreement among annotators, (ii) plausibility is most likely assigned given more abstract words.


For perturbed implausible triples, the picture looks different with only one abstractness combination for which original and majority rating proportions overlap, namely (a, c, c) . For four highly abstract combinations (a, m, a) , (m, a, a) , (m, m, a) , (m, m, a) , a plausible majority is observed. However, in comparison with attested plausible triples, disagreement and uncertainty is much higher with no clear majority for 80% of abstractness combinations. These findings underline the observations for attested plausible triples with (i) implausibility being easier to catch given concrete words, and (ii) plausibility connected to more abstract words.

Abstractness at Event Constituent Level We further examine abstractness at constituent level, i.e., we explore whether abstractness degrees of individual constituents play a role. For this, we again calculate strict majority ratings across triples for each abstractness combination in a binary label setup (cf. 5.2). We focus on triples with a $\geq 70\%$ majority for either plausible or implausible and calculate the proportion of concrete, mid-range, and abstract constituents $\in \{s, v, o\} \in t$.

Results are presented in Table 1. For constituents of triples receiving plausible majority votings, no particular pattern stands out: we find relatively equal shares for all constituents across abstractness levels. For originally implausible triples rated plausible, we observe a slightly higher share of mid-range and abstract constituents. In contrast, abstractness levels seem to play a more important



(a) Attested plausible triples.



(b) Perturbed implausible triples.

Figure 5: Proportion of strict majority ratings ($\geq 70\%$) across abstractness combinations for attested plausible triples (a) and perturbed implausible triples (b). Green bars denote a majority of plausible ratings $\in \{4, 5\}$, pink bars refer to a majority of implausible ratings $\in \{1, 2\}$, and orange bars capture cases of no clear majority.

role for constituents of triples with implausible majority votings. For both originally plausible and implausible triples, percentage shares clearly increase for concrete subjects and objects as compared to triples with plausible majorities. We also observe more abstract verbs, while shares of concrete and mid-range verbs decrease. In addition, a decrease in abstract subjects and objects as well as mid-range subjects can be observed. Regarding verb constituents, the line seems to be clear-cut between verbs as we find an increase in abstract, a decrease in mid-range, and relatively equal shares for concrete verbs.

These examinations suggest that abstractness levels of event constituents are especially important when assessing the absence of plausibility. Generally, events with a majority voting for implausible tend to include more concrete subjects and objects. However, the picture gets more diverse with clear increases in abstract verbs. Interestingly, these observations hold irrespective of the original label.

The exploration of abstractness at event constituents underlines our findings from the previous analysis focusing on abstractness at event level. We again find that the majority of human annotators tend to agree on what is plausible, while implausibility seems to be harder to catch and introduces more disagreement. Moreover, assignment **likelihood of plausibility increases with abstractness** of triple constituents, whereas assignment **likelihood of implausibility increases with concreteness** of triple constituents – no matter the underlying original label.

6 Final Dataset: Aggregations

To foster learning with and from disagreement, we release not only (i) the raw annotator ratings, but also (ii) provide the following standard aggregations to enable various perspectives for interpretation and modelling; for further aggregation options see e.g., Uma et al. (2021b). We account for both multi-class (label $\in \{1, 2, 4, 5\}$) and binary (label either plausible $\in \{4, 5\}$ or implausible $\in \{1, 2\}$) categorizations. The dataset is available at <https://github.com/AnneroseEichel/PAP>.

(a) Strict Majority with Disagreement

Classes are assigned based on a 70% majority for a multi-class or binary setup. In case of no clear majority, a label denoting disagreement is assigned to reflect conflicting perspectives of annotators.

(b) Distribution

To account for fine-grained disagreement and uncertainty, we calculate class distributions for a multi-class or binary setup.

(c) Probabilistic Aggregation

As we work with crowd workers, we also provide probabilistic label aggregations using Multi-Annotator Competence Estimation (MACE)¹² (Hovy et al., 2013). MACE leverages an unsupervised item-response model that learns to identify trustworthy crowd annotators and predicts the correct underlying label. We

¹²<https://github.com/dirkhovy/MACE>

among 1,733 valid triples	originally plausible			originally implausible		
	maj. plausible	maj. implausible	no maj.	maj. plausible	maj. implausible	no maj.
# triples	622	11	229	309	46	516
constituent	constituent proportion (in %)					
concrete subjects	0.106	0.182	0.158	0.100	0.152	0.144
concrete verbs	0.109	0.091	0.162	0.093	0.116	0.172
concrete objects	0.100	0.182	0.070	0.088	0.159	0.060
mid-range subjects	0.115	0.061	0.129	0.118	0.058	0.144
mid-range verbs	0.115	0.091	0.034	0.128	0.072	0.132
mid-range objects	0.111	0.061	0.125	0.111	0.138	0.042
abstract subjects	0.113	0.091	0.148	0.115	0.123	0.146
abstract verbs	0.109	0.152	0.148	0.112	0.145	0.129
abstract objects	0.122	0.091	0.025	0.134	0.036	0.031

Table 1: Overview of constituent analysis focusing on triples with a $\geq 70\%$ majority (maj.) for either plausible or implausible triples (# triples). We present the proportion of concrete, mid-range, and abstract constituents $\in \{s, v, o\} \in t$ for each abstractness level (concrete, mid-range, abstract) and constituent (subject, verb, object), in %. For completeness, we also show constituent proportions for triples with no strict majority (no maj.).

provide both predicted silver labels and class distributions for a multi-class and binary setup.

7 Discussion

We formulated the task of automatically distinguishing *abstract* plausible events from implausible ones as an extension of Wang et al. (2018) who focused specifically on *physical* plausible events. Based on the presented findings, we affirm our hypothesis as to (i) whether plausible and implausible events can be systematically captured on physical and abstract levels by (ii) integrating degrees of abstractness for combinations of words.

We further note differences in collected annotations with assignment likelihood of plausible ratings increasing with abstractness of events’ constituents, while concreteness seems to facilitate the detection of more implausible events. We hypothesize that more concrete words evoke a more stable mental image grounded in the real world. Events like our introductory example *rain-breaks-belly* that represent a violation of quite fixed mental images are thus more often recognized as implausible. In contrast, more abstract words that lack a tangible reference object seem to open up a greater space of potentially plausible interpretations. This possibly invites annotators to cooperate and use their imagination resulting in more plausible ratings for more abstract triples.

Our findings further suggest that it is the recipient who comes up with an interpretation, thus making sense of the seemingly implausible. Moreover, generating fully implausible events is not trivial, which should be taken into account when using automatically generated implausible triples.

Lastly, while events based on *s-v-o* triples or comparably simple constructions have been successfully leveraged for exploring selection preference and thematic fit (Erk et al., 2010; Zhang et al., 2019; Pedinotti et al., 2021), the addition of context exceeding sentences constructed from *s-v-o* triples could potentially resolve present ambiguity and possibly reduce disagreement. We thus encourage future work extending this work by collecting and analyzing plausibility ratings for more complex constructions within broader contexts.

8 Conclusion

We presented a novel dataset for physical and abstract plausibility for events in English. Based on naturally occurring sentences extracted from Wikipedia, we infiltrated degrees of abstractness, and automatically generated perturbed pseudo-implausible events. We annotated a filtered and balanced dataset for plausibility using crowd-sourcing and performed extensive cleaning steps to ensure annotation quality. We provided in-depth analyses to explore the relationship between abstractness and plausibility and examined annotator disagreement. We hope that the presented dataset is used for both analyzing and modelling the notion of plausibility as well as the exploration of closely related tasks such as selectional preference and thematic fit and relevant downstream tasks including common-sense reasoning, NLI, and coreference resolution. Moreover, we make both raw annotations and a range of aggregations publicly available to foster research on disagreement and enable interpretation from various perspectives.

Limitations

In this paper, we present a collection of plausibility ratings for simple sentences in English that are automatically constructed from *s-v-o* triples that are extracted from natural language. We are aware that, for example, events such as *eat-skyscraper* might have a plausible interpretation in a given fictional world. When constructing our dataset, we do not explicitly account for triples which might originate from Wikipedia articles with content where other possible worlds are assumed.

As we conduct a relatively large annotation experiment via AMT crowd-sourcing, we aim to apply post-processing methods minimising the impact of unreliable annotations on our analyses. With more than 500 different final annotators and a very subjective annotation task, we however note the possibility of potentially wrong annotations due to errors, limitations of task instructions, or the interface (Pradhan et al., 2012; Poesio et al., 2019; Uma et al., 2022). This is especially true for the implausible portion of the dataset where no comparison with an attested triple label is possible. Approaches of mitigation could be concentrating on triples with high (im)plausibility ratings or use e.g., probabilistic methods to aggregate labels. We thus provide a dataset version with labels aggregated using MACE (Hovy et al., 2013).

As far as the transfer of the suggested approach of dataset construction to languages other than English is concerned, we call attention to the potential need to adapt the event extraction. Further, abstractness ratings might not readily be available in every language. In addition, AMT annotation for languages other than English potentially requires more time and resources, as annotator population is heavily skewed towards speakers of English.

Ethics Statement

To generate our dataset of events, we use a portion of the English Wikipedia which has been shown to exhibit a range of biases (Olteanu et al., 2019; Schmahl et al., 2020; Falenska and Çetinoğlu, 2021; Sun and Peng, 2021). While our goal is to enable others to explore plausibility on physical and abstract levels as well as sources of potential disagreement, users of this dataset should acknowledge potential biases and should not use to make deployment decisions or rule out failures.

In the context of our annotation task, we collected plausibility ratings from crowd-workers us-

ing Amazon Mechanical Turk between January, 20 and March 7, 2023. Crowd-workers were compensated 0.02\$ per instance. Although we aimed for strict quality control during data collection, we mostly compensated completed hits also when annotations were finally discarded because they did fail a check instance or, sometimes, did not move the slider. To this end, we engaged in email conversations with crowd-workers in case they reached out to clarify issues. We invested time to answer all requests and made our decision-making transparent to the annotators.

Acknowledgements

We are grateful to Prisca Piccirilli, Neele Falk, and the whole SemRel group for helpful suggestions and feedback regarding this work. We would also like to thank the anonymous reviewers for their nuanced comments and constructive suggestions. Annerose Eichel received funding by the Hanns Seidel Foundation’s Talent Program.

References

- Dina Almanea and Massimo Poesio. 2022. [ArMIS - the Arabic misogyny and sexism corpus with annotator subjective disagreements](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2282–2291, Marseille, France. European Language Resources Association.
- Talita Anthonio, Anna Sauer, and Michael Roth. 2022. [Clarifying Implicit and Underspecified Phrases in Instructional Text](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3319–3330, Marseille, France. European Language Resources Association.
- Lora Aroyo and Chris Welty. 2015. Truth is a lie: Crowd truth and the seven myths of human annotation. *AI Magazine*, 36(1):15–24.
- Joris Baan, Wilker Aziz, Barbara Plank, and Raquel Fernandez. 2022. [Stop measuring calibration when humans disagree](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1892–1915, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Lawrence W Barsalou and Katja Wiemer-Hastings. 2005. Situating Abstract Concepts. *Grounding cognition: The role of perception and action in memory, language, and thought*, pages 129–163.
- Valerio Basile, Federico Cabitza, Andrea Campagner, and Michael Fell. 2021a. [Toward a perspectivist turn in ground truthing for predictive computing](#).

- Valerio Basile, Michael Fell, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, Massimo Poesio, and Alexandra Uma. 2021b. [We need to consider disagreement in evaluation](#). In *Proceedings of the 1st Workshop on Benchmarking: Past, Present and Future*, pages 15–21, Online. Association for Computational Linguistics.
- Julia Bettinger, Anna Hättö, Michael Dorna, and Sabine Schulte im Walde. 2020. [A domain-specific dataset of difficulty ratings for German noun compounds in the domains DIY, cooking and automotive](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4359–4367, Marseille, France. European Language Resources Association.
- Marc Brysbaert, Boris New, and Emmanuel Keuleers. 2012. Adding part-of-speech information to the SUBTLEX-US word frequencies. *Behavior research methods*, 44(4):991–997.
- Marc Brysbaert, Amy Beth Warriner, and Victor Kuperman. 2014. Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior research methods*, 46(3):904–911.
- Jean Charbonnier and Christian Wartena. 2019. [Predicting word concreteness and imagery](#). In *Proceedings of the 13th International Conference on Computational Semantics - Long Papers*, pages 176–187, Gothenburg, Sweden. Association for Computational Linguistics.
- Emmanuele Chersoni, Philippe Blache, and Alessandro Lenci. 2016. [Towards a distributional model of semantic complexity](#). In *Proceedings of the Workshop on Computational Linguistics for Linguistic Complexity (CL4LC)*, pages 12–22, Osaka, Japan. The COLING 2016 Organizing Committee.
- Guy Dove. 2009. Beyond perceptual symbols: A call for representational pluralism. *Cognition*, 110(3):412–431.
- Ali Emami, Ian Porada, Alexandra Olteanu, Kaheer Suleman, Adam Trischler, and Jackie Chi Kit Cheung. 2021. [ADEPT: An Adjective-Dependent Plausibility Task](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7117–7128, Online.
- Katrin Erk, Sebastian Padó, and Ulrike Padó. 2010. [A flexible, corpus-driven model of regular and inverse selectional preferences](#). *Computational Linguistics*, 36(4):723–763.
- Agnieszka Falenska and Özlem Çetinoğlu. 2021. [Assessing gender bias in Wikipedia: Inequalities in article titles](#). In *Proceedings of the 3rd Workshop on Gender Bias in Natural Language Processing*, pages 75–85, Online. Association for Computational Linguistics.
- Tommaso Fornaciari, Alexandra Uma, Silviu Paun, Barbara Plank, Dirk Hovy, and Massimo Poesio. 2021. [Beyond black & white: Leveraging annotator disagreement via soft-label multi-task learning](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2591–2597, Online. Association for Computational Linguistics.
- Diego Frassinelli and Sabine Schulte im Walde. 2019. [Distributional interaction of concreteness and abstractness in verb–noun subcategorisation](#). In *Proceedings of the 13th International Conference on Computational Semantics - Short Papers*, pages 38–43, Gothenburg, Sweden. Association for Computational Linguistics.
- Kevin Gimpel, Nathan Schneider, Brendan O’Connor, Dipanjan Das, Daniel Mills, Jacob Eisenstein, Michael Heilman, Dani Yogatama, Jeffrey Flanigan, and Noah A. Smith. 2011. [Part-of-speech tagging for Twitter: Annotation, features, and experiments](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 42–47, Portland, Oregon, USA. Association for Computational Linguistics.
- Jonathan Gordon and Benjamin Van Durme. 2013. [Reporting bias and knowledge acquisition](#). In *Proceedings of the 2013 Workshop on Automated Knowledge Base Construction*, page 25–30, New York, NY, USA. Association for Computing Machinery.
- Dirk Hovy, Taylor Berg-Kirkpatrick, Ashish Vaswani, and Eduard Hovy. 2013. [Learning whom to trust with MACE](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1120–1130, Atlanta, Georgia. Association for Computational Linguistics.
- Dirk Hovy, Barbara Plank, and Anders Søgaard. 2014. [Experiments with crowdsourced re-annotation of a POS tagging data set](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 377–382, Baltimore, Maryland. Association for Computational Linguistics.
- Paul Jaccard. 1902. [Lois de distribution florale dans la zone alpine](#). *Bulletin de la Société vaudoise des sciences naturelles*, 38:69–130.
- Emily Jamison and Iryna Gurevych. 2015. [Noise or additional information? Leveraging crowdsource annotation item agreement for natural language tasks](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 291–297, Lisbon, Portugal. Association for Computational Linguistics.
- Kian Kenyon-Dean, Eisha Ahmed, Scott Fujimoto, Jeremy Georges-Filteau, Christopher Glasz, Barleen Kaur, Auguste Lalande, Shruti Bhandari, Robert

- Belfer, Nirmal Kanagasabai, Roman Sarrazingen-dron, Rohit Verma, and Derek Ruths. 2018. [Sentiment analysis: It's complicated!](#) In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1886–1895, New Orleans, Louisiana. Association for Computational Linguistics.
- Elisa Leonardelli, Stefano Menini, Alessio Palmero Aprosio, Marco Guerini, and Sara Tonelli. 2021. [Agreeing to disagree: Annotating offensive language datasets with annotators' disagreement.](#) In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10528–10539, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Eleni Metheniti, Tim Van de Cruys, and Nabil Hathout. 2020. [How relevant are selectional preferences for transformer-based language models?](#) In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1266–1278, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Daniela Naumann, Diego Frassinelli, and Sabine Schulte im Walde. 2018. [Quantitative semantic variation in the contexts of concrete and abstract words.](#) In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 76–85, New Orleans, Louisiana. Association for Computational Linguistics.
- Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Emre Kıcıman. 2019. [Social data: Biases, methodological pitfalls, and ethical boundaries.](#) *Frontiers in Big Data*, 2.
- Ellie Pavlick and Tom Kwiatkowski. 2019. [Inherent disagreements in human textual inferences.](#) *Transactions of the Association for Computational Linguistics*, 7:677–694.
- Paolo Pedinotti, Giulia Rambelli, Emmanuele Chersoni, Enrico Santus, Alessandro Lenci, and Philippe Blache. 2021. [Did the cat drink the coffee? Challenging transformers with generalized event knowledge.](#) In *Proceedings of *SEM 2021: The Tenth Joint Conference on Lexical and Computational Semantics*, pages 1–11, Online. Association for Computational Linguistics.
- Barbara Plank, Dirk Hovy, and Anders Søgaard. 2014. [Linguistically debatable or just plain wrong?](#) In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 507–511, Baltimore, Maryland. Association for Computational Linguistics.
- Massimo Poesio, Jon Chamberlain, Silviu Paun, Juntao Yu, Alexandra Uma, and Udo Kruschwitz. 2019. [A crowdsourced corpus of multiple judgments and disagreement on anaphoric interpretation.](#) In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1778–1789, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ian Porada, Kaheer Suleman, and Jackie Chi Kit Cheung. 2019. [Can a gorilla ride a camel? Learning semantic plausibility from text.](#) In *Proceedings of the First Workshop on Commonsense Inference in Natural Language Processing*, pages 123–129, Hong Kong, China.
- Ian Porada, Kaheer Suleman, Adam Trischler, and Jackie Chi Kit Cheung. 2021. [Modeling event plausibility with consistent conceptual abstraction.](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1732–1743, Online.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. [CoNLL-2012 Shared Task: Modeling multilingual unrestricted coreference in OntoNotes.](#) In *Joint Conference on EMNLP and CoNLL - Shared Task*, pages 1–40, Jeju Island, Korea. Association for Computational Linguistics.
- Valentina Pyatkin, Shoval Sadde, Aynat Rubinstein, Paul Portner, and Reut Tsarfaty. 2021. [The possible, the plausible, and the desirable: Event-based modality detection for language processing.](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 953–965, Online. Association for Computational Linguistics.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. Stanza: A Python natural language processing toolkit for many human languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*.
- Radim Řehůřek and Petr Sojka. 2010. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pages 45–50, Valletta, Malta. ELRA. <http://is.muni.cz/publication/884893/en>.
- Philip Stuart Resnik. 1993. *Selection and information: A class-based approach to lexical relationships*. Ph.D. thesis, University of Pennsylvania.
- Filipe Rodrigues, Mariana Lourenço, Bernardete Ribeiro, and Francisco C. Pereira. 2017. [Learning supervised topic models for classification and regression from crowds.](#) *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(12):2409–2422.

- Stephen Roller, Sabine Schulte im Walde, and Silke Scheible. 2013. [The \(un\)expected effects of applying standard cleansing models to human ratings on compositionality](#). In *Proceedings of the 9th Workshop on Multiword Expressions*, pages 32–41, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Asad Sayeed, Clayton Greenberg, and Vera Demberg. 2016. [Thematic fit evaluation: an aspect of selectional preferences](#). In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*, pages 99–105, Berlin, Germany. Association for Computational Linguistics.
- Roland Schäfer. 2015. [Processing and querying large web corpora with the COW14 architecture](#). *Proceedings of the 3rd Workshop on Challenges in the Management of Large Corpora*, Lancaster, 20 July 2015, pages 28 – 34, Mannheim.
- Katja Geertruida Schmahl, Tom Julian Viering, Stavros Makrodimitris, Arman Naseri Jahfari, David Tax, and Marco Loog. 2020. [Is Wikipedia succeeding in reducing gender bias? Assessing changes in gender bias in Wikipedia using word embeddings](#). In *Proceedings of the Fourth Workshop on Natural Language Processing and Computational Social Science*, pages 94–103, Online. Association for Computational Linguistics.
- Sabine Schulte im Walde and Diego Frassinelli. 2022. [Distributional Measures of Abstraction](#). *Frontiers in Artificial Intelligence: Language and Computation* 4:796756. Alessandro Lenci and Sebastian Pado (topic editors): "Perspectives for Natural Language Processing between AI, Linguistics and Cognitive Science".
- Vered Shwartz and Yejin Choi. 2020. [Do neural language models overcome reporting bias?](#) In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6863–6870, Barcelona, Spain (Online).
- Jiao Sun and Nanyun Peng. 2021. [Men are elected, women are married: Events gender bias on Wikipedia](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 350–360, Online. Association for Computational Linguistics.
- Tarun Tater, Diego Frassinelli, and Sabine Schulte im Walde. 2022. [Concreteness vs. abstractness: A selectional preference perspective](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing: Student Research Workshop*, pages 92–98, Online. Association for Computational Linguistics.
- Alexandra Uma, Dina Almanea, and Massimo Poesio. 2022. Scaling and Disagreements: Bias, Noise, and Ambiguity. *Frontiers in Artificial Intelligence*, 5.
- Alexandra Uma, Tommaso Fornaciari, Anca Dumitrache, Tristan Miller, Jon Chamberlain, Barbara Plank, Edwin Simpson, and Massimo Poesio. 2021a. [SemEval-2021 Task 12: Learning with disagreements](#). In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 338–347, Online. Association for Computational Linguistics.
- Alexandra Uma, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, and Massimo Poesio. 2021b. Learning from disagreement: A survey. *Journal of Artificial Intelligence Research*, 72:1385–1470.
- Wessel Van Dam, Shirley-Ann Rueschemeyer, Oliver Lindemann, and Harold Bekkering. 2010. [Context effects in embodied lexical-semantic processing](#). *Frontiers in Psychology*, 1:150.
- Tim Van de Cruys. 2014. [A neural network approach to selectional preference acquisition](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 26–35, Doha, Qatar. Association for Computational Linguistics.
- Su Wang, Greg Durrett, and Katrin Erk. 2018. [Modeling Semantic Plausibility by Injecting World Knowledge](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 303–308, New Orleans, Louisiana.
- Yorick Wilks. 1975. [A preferential, pattern-seeking, semantics for natural language inference](#). *Artificial Intelligence*, 6(1):53–74.
- Jie Yin, Nalin Narang, Paul Thomas, and Cecile Paris. 2012. [Unifying local and global agreement and disagreement classification in online debates](#). In *Proceedings of the 3rd Workshop in Computational Approaches to Subjectivity and Sentiment Analysis*, pages 61–69, Jeju, Korea. Association for Computational Linguistics.
- Hongming Zhang, Hantian Ding, and Yangqiu Song. 2019. [SP-10K: A large-scale evaluation set for selectional preference acquisition](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 722–731, Florence, Italy. Association for Computational Linguistics.

A Dataset Construction

A.1 Filtering the Brysbaert Norms

To assign abstractness scores to all nouns and verbs in a given event triples, we draw on the concreteness ratings for approximately 40,000 English words devised by [Brysbaert et al. \(2014\)](#). The Brysbaert norms were collected in an out-of-context setting and without providing information about the part-of-speech (POS). POS tags were added

in a post-processing step, utilizing the SUBTLEX-US corpus (Brysbaert et al., 2012). To account for this, we follow Schulte im Walde and Frassinelli (2022) and Tater et al. (2022) in adding the most frequent POS tag associated with each target word based on the English web corpus ENCOW16AX (Schäfer, 2015). We then filter for noun and verb target words where the POS tag provided by (Brysbaert et al., 2014) and the POS tag extracted using the ENCOW16AX correspond to each other. We filter out all words with a frequency below 10K to remove infrequent words. This way, we obtain a collection of 12,880 noun and 2,522 verb targets.

A.2 Triple Binning and Distributions

The distribution of all naturally occurring triples for each triple combination $\in \{(a, a, a), \dots (c, c, c)\}$ is presented in Fig. 6 (a). While triple numbers accumulate on the extremes highly abstract and highly concrete, the number drops for triples consisting of mid-range constituents. Mixed triple combinations (a, m, c) and (c, m, a) yield minimum numbers of triples as well as triples with highly concrete or abstract subjects and verbs (a, a, c) and (c, c, a) .

Similarly, the distribution of all automatically generated pseudo-implausible triples for each triple combination is shown in Fig. 6 (b). Note that a substantially higher number of valid implausible triples is extracted using the binning process with minimum numbers achieved for mostly medium-range abstractness.

B Human Annotation

B.1 HIT Interface

Fig. 7 shows a full example of the HIT interface as presented to HIT workers.

B.2 Check Instances

We list check instances in Table 2. In a post-processing step, we exclude three implausible check instances, e.g., *water cuts ball*, which might be interpreted as plausible in the context of high-pressure water systems which might be able to cut a ball (marked in italics). We use the check instances mainly after the the annotations process to increase annotation quality by filtering out all submissions where annotators failed a valid check instance.

B.3 Annotation Post-Processing

We show an overview of dataset statistics at each post-processing step in Table 3. Specifically, we

plausible	implausible
grandmother drinks tea	grandmother drinks stone
child eats banana	child eats dream
baker bakes cake	<i>baker bakes air</i>
kid plays game	sun beats banana
rabbit eats carrot	baby eats storm
man builds house	man breaks air
man opens window	ant opens window
teenager drinks coke	<i>sun breaks door</i>
woman drives car	woman drinks bridge
player throws ball	<i>water cuts ball</i>

Table 2: Plausible and implausible check instances. Instances marked in italics are filtered out in a post-processing step due to possible plausible interpretations.

present changes in number of ratings, validated annotators, and number of triples with >8 ratings across annotation post-processing. Post-processing methods are applied in the order listed. Results in a given row correspond to dataset statistics having applied a given step.

Soft Jaccard Coefficient We estimate Inter-Annotator Agreement (IAA) by calculating the Jaccard Coefficient for all pairwise annotator combinations

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|}$$

where the intersection of A and B captures all cases where annotators agree that a triple is either plausible (ratings $\in \{4, 5\}$) or implausible (ratings $\in \{1, 2\}$), and the union of A and B denotes all cases where both annotators provided a rating for the same sentence regardless of their agreement. As we are not enforcing strict rating agreement, we refer to this way of calculating IAA as *soft* Jaccard Coefficient.

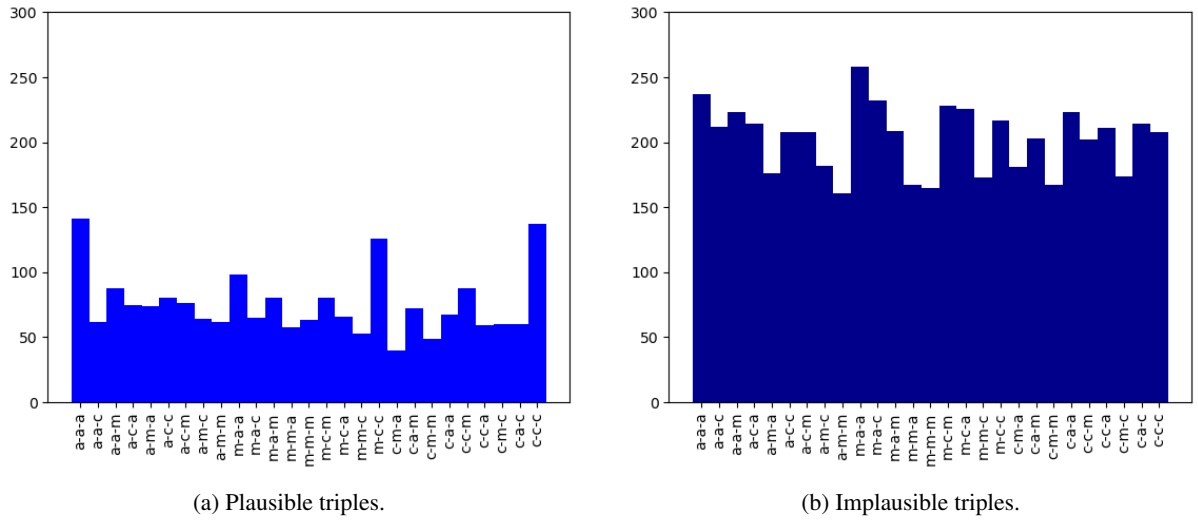


Figure 6: Distribution of attested plausible (left) and perturbed implausible (right) triples per combination.

Instructions

Task: Please read the sentence carefully. Then decide whether the sentence is plausible or implausible and move the slider in the corresponding direction.

Try not to think too much and trust your gut feeling!

- **Note that you have to move the slider in one direction.** Ratings where the slider is set to the middle will be rejected.
- Please answer all 25 HITs.
- This experiment is for **English native speakers** only.

Your Task:

Is the following sentence plausible or not?

"paper introduces innovation"

Move the slider in the corresponding direction.

implausible plausible

Any comments?

Figure 7: HIT interface including task instruction and requirements for successful answer submission (ratings where the slider is set to the middle can be rejected).

	# annotators		# ratings		# triples	
	plausible	implausible	plausible	implausible	total	total
Raw (without check instances)	883	879	11,250	11,343	22,593	2,160
Failed checks/default submission	878	872	10,649	10,668	21,317	2,148
>75% disagreement orig. label	832	838	9,849	10,046	19,895	2,081
<10 ratings submitted	478	479	8,373	8,502	16,875	1,840
AMT approv. rate <80%, incl. 0%	468	471	8,269	8,333	16,602	1,819
Pairwise Jaccard Index <0.4	452	452	7,726	7,845	15,571	1,733

Table 3: Overview of changes in number of ratings, validated annotators, and number of triples with >8 ratings across annotation post-processing. Post-processing methods are applied in the order listed. Results in a given row correspond to dataset statistics having applied a given step, e.g., filtering out submission from annotators who failed check instances as well as all submissions where annotators submitted a default rating of 3 results in the number of 21,317 valid ratings, including both ratings for plausible and implausible triples.