

Effects of entity relatedness and definiteness on bridging inferences

Mandy Simons, Carnegie Mellon University

Hannah Rohde, University of Edinburgh

An interpreter encountering an NP in discourse must decide whether its referent is part of a situation already constructed in their mental model (bridged interpretation), or is a new, unrelated entity. In (1), an entity *the living room* is introduced in a context sentence; subsequent NPs can be related to this situation or interpreted to refer to an unrelated entity.

- (1) a. Jane was standing in the living room. The window...
b. Jane was standing in the living room. The congresswoman...
c. Jane was standing in the living room. A congresswoman...

A bridged interpretation for (1a) is easily inferable, with the window understood as a window in the living room (e.g., *The window was open and Jane could feel a breeze*) but a non-bridged reading may also be coherent (e.g., *The window of a car that drove by was rolled down and she could hear music blaring*). Work in formal pragmatics identifies two factors that are claimed to contribute to bridging inferences: entity relatedness (Asher & Lascarides 1998; Prince 1992) and definiteness (Clark 1975; Roberts 2003). For (1b), a bridged interpretation may be disfavored given the atypicality of a congresswoman in the living room, and for (1c) the indefinite determiner may undermine the referent's givenness (no bridge: *The/A congresswoman announced that the state was going into lockdown*). Experimental work on NP processing has found early robust effects of entity relatedness and also some influence of definiteness but only in later measures (N400 vs P600 in ERP; Schumacher 2009). This work, however, doesn't establish how to determine whether a comprehender has indeed established a bridged interpretation. Such interpretations are often assumed to arise (as in Clark's examples or Schumacher's materials), but to test what factors support bridging inferences, we need a clear metric of whether comprehenders indeed infer a bridged interpretation. Here we present three experiments that manipulate entity relatedness and definiteness while testing for the presence of bridging. The results show that entity relatedness affects the interpretation of an NP (offline & online); definiteness alone does not influence the interpretation but high-related definites favor bridged interpretations (online).

Experiment 1. Our goal is to test which properties of an NP encourage a bridged interpretation by assessing how participants treat that NP in a story continuation task. Participants (N=54, mturk) wrote story continuations for 40 targets, 40 fillers. Target items described a context followed by the potentially bridgeable NP (high vs low related; def vs indef), as in (2).

- (2) [high, def] Ian likes to work at a large desk. The chair _____
[high, indef] Ian likes to work at a large desk. A chair _____
[low, def] Hilda created a nice arrangement of fruit. The chair _____
[low, indef] Hilda created a nice arrangement of fruit. A chair _____

The annotation process showed many cases of unambiguous bridging (*Ian likes to work at a large desk. The chair ... fits nicely underneath*) and cases of likely non-bridging (*Hilda created a nice arrangement of fruit. A chair... was on the porch*), but many were equally coherent assuming a bridged or a non-bridged interpretation. Relying only on the annotators' intuitions risked a circular treatment in which the coding would reflect the annotators' own sensitivity to the manipulated factors rather than their effect on participants. For example, similar continuations might be treated as bridges for high-related entities (*Ian likes to work at a large desk. The chair... is very comfy*) and non-bridges for low-related entities (*Hilda created a nice arrangement of fruit. The chair... is very nice*). Experiment 1 was thus inconclusive, but we used the resulting continuations to create a different paradigm in Experiment 2.

Experiment 2. To assess the potential bridges in participants' continuations in Expt 1, we asked new participants (N=55, mturk) to complete dialogues that probe whether the continuations generated in Expt 1 are interpreted by the new participants as bridged or not, as in (3,4). We used the Expt 1 data with 4 randomly selected continuations per NP per condition.

- (3) Speaker A: Ian likes to work at a large desk. The chair leans back and was quite expensive.
Speaker B: Wait, sorry, I wasn't listening. Which chair are you talking about?
Speaker A: _____
- (4) Speaker A: Hilda created a nice arrangement of fruit. The chair had dust on it.
Speaker B: Wait, sorry, I wasn't listening. Which chair are you talking about?
Speaker A: _____

Our assumption was that a bridged interpretation would yield Speaker A replies that repeated content from the context sentence. To illustrate with two sample replies, participants wrote *The one at the big desk Ian likes to work at* for (3) and *The chair that had dust on it* for (4), where the former indicates a bridged reading and the latter does not. We used a string similarity word-overlap metric whose scores we modelled with linear mixed-effect regressions. In using a continuous measure to identify bridging, we acknowledge that bridging may be a matter of degree and not a binary feature. For context~reply similarity, we found higher scores in the high-related than low-related condition ($p < .001$) and no effect of definiteness or interaction, suggesting that bridging primarily reflects entity relatedness. For continuation~reply similarity, we found higher scores for the low-related than high-related condition ($p < .001$); we also found a relatedness X definiteness interaction ($p < .001$), suggesting a pattern whereby more cases of non-bridging emerged in the low-related condition, particularly for indefinites in that condition.

Experiment 3. In a self-paced reading paradigm, we assessed RTs at the point in a sentence where a potential bridge is cancelled, with the prediction that factors that support a bridging inference will increase the processing difficulty if that bridge must be cancelled. Participants (N=100, prolific) read passages (40 target, 24 filler), with target items that consisted of a context sentence followed by a continuation with multi-word chunks for the determiner-noun, the start of a relative clause (RC), bridge-incompatible content in the RC, and two or more spillover regions.

- (5) Context sentence: Jane was in the living room.
[high, def] The window that was in **her dream** suddenly came to mind.
[high, indef] A window that was in **her dream** suddenly came to mind.
[low, def] The knife that was in **her dream** suddenly came to mind.
[low, indef] A knife that was in **her dream** suddenly came to mind.

For RTs at the critical bridge-incompatible region (bold in (5)), a linear mixed effects model showed a main effect of relatedness ($p < .05$) with slower RTs for high-related nouns. There was also a relatedness X definiteness interaction ($p < .05$), whereby the high-related definite condition had the slowest RTs. We take these results to show that entity relatedness is a core component for bridged interpretations and that definiteness, while it alone does not trigger bridging (i.e., low-related definites didn't give rise to the slowdown that would be evidence of bridging), acts together with high relatedness to make bridged interpretations more likely.

Overall, these results contribute to theoretical models of bridging and more broadly to a growing literature challenging analyses of definiteness as an unambiguous signal of givenness or contextual uniqueness.

Asher & Lascarides 1998. *Jrnl of Semantics*. **Clark 1975.** In Schank & Nash-Webber's *Theoretical issues in natural language processing*. **Prince 1992.** In Mann & Thompson's *Discourse description: Diverse linguistic analyses of a fund-raising text*. **Roberts 2003.** *Linguistics & Philosophy*. **Schumacher 2009.** In Lalitha Devi et al.'s *Anaphora processing*.