Running head: Keeping track of co-presence

What is retained about common ground?
Distinct effects of linguistic and visual co-presence

Alexia Galati ^a and Susan E. Brennan ^b

^a Department of Psychological Science, University of North Carolina at Charlotte,
Charlotte, NC, USA

^b Department of Psychology, Stony Brook University, Stony Brook, NY, USA

Accepted in Cognition (06/08/2021)

Authors' note

This material is based upon work supported by NSF under Grant # ITR-0325188. We are thankful to Jenny Lee, Billy Galbavy, and Gennadiy Ryklin for assistance with coding, and to Arty Samuel, Amanda Stent, Hoi-Chung Leung, and Richard Gerrig for useful discussions. We also thank the two anonymous reviewers for their constructive comments.

Correspondence concerning this article should be addressed to Alexia Galati, Department of Psychological Science, University of North Carolina at Charlotte. Email: agalati@uncc.edu

Abstract

Common ground can be mutually established between conversational partners in several ways. We examined whether the modality (visual or linguistic) with which speakers share information with their conversational partners results in memory traces that affect subsequent references addressed to a particular partner. In 32 triads, directors arranged a set of tangram cards with one matcher and then with another, but in different modalities, sharing some cards only linguistically (by describing cards the matcher couldn't see), some only visually (by silently showing them), some both linguistically and visually, and others not at all. Then directors arranged the cards again in separate rounds with each matcher. The modality with which they previously established common ground about a particular card with a particular matcher (e.g., linguistically with one partner and visually with the other) affected subsequent referring: References to cards previously shared only visually included more idea units, words, and reconceptualizations than those shared only linguistically, which in turn included more idea units, words, and reconceptualizations than those shared both linguistically and visually. Moreover, speakers were able to tailor references to the same card appropriately to the distinct modality shared with each addressee. Such gradient, partner-specific adaptation during re-referring suggests that memory encodes rich-enough representations of multimodal shared experiences to effectively cue relevant constraints about the perceptual conditions under which speakers and addressees establish common ground.

Keywords: audience design, common ground, communication modality, co-presence, partner-specific adaptation, speech planning

Introduction

When interacting with others across space and time, people share information through various modalities; they adapt the strategies they use to express information depending on whether they are able to hear their conversational partners, see them, or physically share a task environment with them (Clark & Brennan, 1991). According to Clark and Marshall (1981), speakers can rely on heuristics from shared conversations, from perceptually shared settings, or from being members of the same community, as a means of making tractable what information is in “common ground” (cf., Stalnaker, 1974). In this view, common ground can be characterized in terms of the conditions of co-presence under which information is shared: namely, *linguistic co-presence* (sharing information through spoken utterances), and *visual co-presence* (sharing information through the physical environment) (Clark & Marshall, 1981). There is evidence that linguistic and visual co-presence each shape how speakers refer to entities (e.g., Brennan, 2005; Clark & Krych, 2004; Gergle, Kraut, & Fussell, 2004), that speakers can keep track of what they've discussed with whom (e.g., Brennan & Clark, 1996; Galati & Brennan, 2010; Horton & Gerrig, 2005b; Yoon & Brown-Schmidt, 2018), and that speakers can design referring expressions understandable to two addressees (with differing knowledge and individual perspectives) in the same multiparty conversation (Yoon & Brown-Schmidt, 2019). However, it is still largely unknown whether speakers can keep track of the co-presence conditions (or perceptual modalities) under which common ground was established with different conversational partners, and appropriately adapt the referring expressions they subsequently address to each partner.

Consider the following scenario: Three friends—Anna, Beth, and Camila—are visiting a museum. The three of them wander around, intercepting each other in different wings of the

museum to look at and sometimes discuss the paintings together. Some paintings Anna views and talks about with Beth, others with Camila; such paintings are shared through both linguistic and visual co-presence between Anna and each of her friends. Some other paintings Anna and Beth view together without discussing them, others Anna and Camila view together without discussion; such paintings are shared just through visual co-presence. When they reunite in another hallway, Anna may also describe to Beth or to Camila some notable artwork she's seen in another wing of the museum that they haven't found yet; such paintings are shared with just linguistic co-presence.

Imagine that at a later point in time, Anna and Beth reminisce about their visit to the museum. Would it be easier for Anna to determine that she indeed shared a particular painting with Beth if it had been one that they had viewed and talked about together, compared to one that she described to Beth in passing, or one that they had viewed together silently? If so, Anna may be more likely to mark information shared both linguistically and visually as being in her common ground with Beth; for instance, she may be more likely to use a definite expression with succinct description (e.g., "I liked the Blue Rider painting"). In contrast, when describing a painting they shared only visually, she may introduce it more tentatively and with more details (e.g., "I liked that kind of Miro-looking, Bauhaus one by Kandinsky").

Much of the research on how speakers assess and use common ground has focused on the time course with which such information impacts language processing. Some researchers have proposed that speakers take into consideration partner-specific information right from the early moments of processing (e.g., Brennan & Hanna, 2009; Hanna, Tanenhaus, Trueswell, 2003; Hanna & Tanenhaus, 2004; Nadig & Sedivy, 2002; Metzing & Brennan, 2003), whereas others have instead proposed that speakers initially consider only egocentric information, with partner-

specific adaptation having its impact later in the form of repairs (e.g., Horton & Keysar, 1996; Keysar, Barr, & Horton, 1998; Keysar, Barr, Balin, & Paek, 1998; Kronmüller & Barr, 2007). Beyond this debate about the time course of accessing partner-specific information, few studies have examined the types of partner-specific distinctions that speakers can retain about shared experiences.

Despite classic evidence that memory encodes mainly gist rather than details (Sachs, 1967), there is evidence that speakers can retain certain kinds of associations between conversational partners and discourse-specific, contextual, or environmental information. For instance, speakers can retain associations between specific conversational partners and the categories of shared information discussed with them, such as “I have talked to Beth about fish and birds and to Camila about turtles and dogs” (e.g., Horton & Gerrig, 2005b, 2002, Brown-Schmidt, Yoon, & Ryskin, 2015). Addressees can also associate a particular expression (e.g., “shiny cylinder”) that they have used repeatedly in a conversation with a particular speaker; if that speaker inexplicably uses a new expression (e.g., “silver pipe”), the addressees experience a delay in resolving the referent, although *not* when a new speaker uses that same new expression (Metzing & Brennan, 2003; Matthews et al., 2010). Even in situations in which people are co-present but lack an explicit intent to communicate, they may be able to use their co-present partner’s identity to cue both category and lexical associations (Horton, 2007; see, however, Brown-Schmidt & Horton, 2014). Speakers can also keep track of whether a referring expression used during conversational interaction was actually grounded with a partner (through a completed acceptance from the partner; see Clark & Schaefer, 1989) or else interrupted, and therefore not reliably established as being in common ground (Brown-Schmidt, 2009). Finally, when the conversational partner’s spatial viewpoint is known in advance, speakers can represent

that viewpoint in memory, as evidenced by performance on memory tests administered prior to describing that spatial information (Galati et al., 2013; Galati & Avraamides, 2015). Similarly, evidence from eye-movements during a collaborative visuospatial task suggests that listeners integrate information about the partner's spatial viewpoint when interpreting spatial instructions, even when doing so taxes cognitive resources (e.g., when partners are spatially misaligned; Ryskin et al., 2014; Ryskin, Wang, & Brown-Schmidt, 2016).

These findings are consistent with a memory-based account for how partner-specific adaptation unfolds in dialogue (Horton & Gerrig, 2005a, 2005b; 2016). According to Horton and Gerrig, information about conversational partners can be retained in episodic traces and is accessed through ordinary memory-dependent processes, such as resonance with combinations of cues in working memory. On this view (see also Metzing & Brennan, 2003), the identity of a conversational partner can serve as a contextual cue for an episode in which partners were perceptually co-present and influence the access of relevant information, probing memory like any other configuration of cues. Failure to make an appropriate adjustment can be explained by difficulty in accessing shared information or making relevant distinctions about the conversational partner's needs. For example, delayed or infelicitous utterances or interpretations can occur when executive functioning or attentional resources are overloaded (Brown-Schmidt, 2009; Epley, Morewedge & Keysar, 2004), when the partner's informational needs require complex inferences that have not yet been made (Gerrig, Brennan, & Ohaeri, 2000), when relevant information isn't available early enough (Kraljic & Brennan, 2005), when there is time pressure (Horton & Keysar, 1996), or when partner-specific associations are weak (Horton & Gerrig, 2005a). In a related proposal, Brown-Schmidt, Yoon, and Ryskin (2015) argue that a conversational partner serves as a contextual memory cue for retrieving jointly experienced

information by providing non-arbitrary and meaningful associations with information from the shared experience. An implication of this proposal is that the effect of the conversational partner on perspective-taking can depend on details such as the type of role the partner had in the interaction (e.g., instructing vs. informing; Yoon, Koh, & Brown-Schmidt, 2012). Insofar as interactions are shaped by the perceptual conditions of co-presence, we ask here whether speakers can retain these meaningful partner-specific associations and use them effectively during subsequent referring.

Keeping track of the conditions of co-presence is relevant to communication because these conditions influence speakers' strategies and resources for *grounding*—the process of seeking and providing evidence about mutual understanding (Brennan, 2005; Brown-Schmidt, 2009; Clark & Brennan, 1991; Clark & Schaefer, 1989; Clark & Wilkes-Gibbs, 1986). For instance, when one partner has visual evidence about what the other understands during a collaborative spatial task, pairs go through a shorter process of verbally ratifying that they mutually understood each other than when visual evidence is lacking (Brennan, 2005; Clark & Krych, 2004). In tasks that are dynamic (e.g., with objects that are changing and hard to describe), visual co-presence helps partners coordinate more easily and achieve better task performance because they are better able to understand the momentary state of the task (Gergle, Kraut, & Fussell, 2004). Speakers can also keep track of the specific manner of their actions that their partner observed. In a computerized Tower of Hanoi task, speakers kept track of whether their partner observed them lifting vs. dragging disks, and adapted their gestures in subsequent explanations of a version of the task that involved new manner information (Hilliard & Cook, 2016). Consistent with Horton and Gerrig (2005a, 2005b; 2016), these findings suggest that

speakers can retain in episodic traces how they have shared and grounded information with a partner.

Given the cognitive demands on working memory and executive functioning that speakers face under the timing or other pressures of dialogue, we have previously proposed that in order to tailor an utterance to a partner's knowledge or needs (a process known as "audience design"), the representation of those needs must be relatively simple and available (Galati & Brennan, 2010, 2014; Brennan, Galati, & Kuhlen, 2010; Brennan & Hanna, 2009). This proposal is consistent with findings that speakers readily adapt to simple, often binary distinctions such as: whether the partner has heard a story before or not (Galati & Brennan, 2010, 2014), whether a particular category of items has been discussed with the partner or not (Horton & Gerrig, 2005b), whether the partner has a picture of what is being described or not (Lockridge & Brennan, 2002), whether the partner sees visual evidence about the state of a collaborative task or not (Brennan, 2005), or whether the partner has referred to a particular object before or not (Metzing & Brennan, 2003). Representing the partner's knowledge or informational needs as simple distinctions that can be easily cued makes tractable the computational challenge that speakers would otherwise face when accessing or assessing those informational needs (see also Butterfill & Apperly, 2013, Shintel & Keysar, 2009, and Duran & Dale, 2014, for other minimalist accounts of tracking partner-specific information).

Despite evidence of speakers' ability to track simple, binary distinctions about co-presence, it is still an open question whether and how they can track multiple distinctions about the informational needs of multiple partners. As we have suggested with our opening example, many naturalistic conversational contexts do not map neatly to a simple, binary distinction about common ground. Indeed, memory for the *source* of information—that is, the conditions under

which a memory was formed—depends on a number of factors. People’s ability to correctly identify or monitor the source of information is shaped by the quality of the information (e.g., perceptual, contextual, affective characteristics, and even information about cognitive operations), the efficacy of their decision processes, and the criteria they set when evaluating these memory records (Johnson, Hashtroudi, & Lindsay, 1993). The process of identifying the source of information involves, among other aspects, assigning weights to various relevant dimensions to set criteria for source assessment (e.g., weighing what someone showed you as a stronger cue than what someone told you, or vice-versa). To the extent that speakers can track different sources of perceptual information from shared multimodal experiences, they should adapt their interpretations or speech plans.

Accounts for modeling perceptual co-presence in audience design

In the present work, we examine whether two different sources of perceptual information—linguistic and visual co-presence—can be tracked in the context of two different addressees, resulting in distinctions among the forms of subsequent referring expressions. Going back to the scenario of Anna, Beth, and Camila’s visit to the museum, more than one co-presence heuristic (or perceptual constraint) is at work: not only whether information was previously discussed with a particular companion (as tested in Galati & Brennan, 2010), but also whether it was shared visually, and with which companion. Thus, the relevant constraints are linguistic co-presence (whether paintings were jointly discussed), visual co-presence (whether paintings were jointly viewed), and partner identity (with which partner this experience was shared). Consider the moment when Anna is reminiscing with Beth about a painting that she (Anna) had described to Beth in the hallway (i.e., a painting shared only through linguistic co-presence). Would she use a different description as when describing a painting she and Beth had

both viewed and discussed together (i.e., a painting shared through both linguistic and visual co-presence)? What about when describing a painting that they both viewed quietly together: one shared by visual co-presence only? We ask whether multimodally shared items are tracked better than items shared through one modality only, reflecting an interaction of the two co-presence constraints. We also ask whether, among unimodally shared items, having shared an item through speech makes it more salient than having shared it visually, reflecting the relative weighing of the two constraints for the task at hand.

We consider two contrasting accounts that can be used (or extended) to characterize the effects of different information sources on partner-specific processing: constraint-based models (Tanenhaus & Trueswell, 1995; McRae, Spivey-Knowlton, & Tanenhaus, 1998) and the interactive alignment model of Pickering and Garrod (2004). Constraint-based models of language processing represent different sources of information as constraints that combine probabilistically to shape the interpretation and production of utterances (e.g., Brown-Schmidt & Hanna, 2011; Tanenhaus & Trueswell, 1995; McRae, Spivey-Knowlton, & Tanenhaus, 1998). As with other sources of information such as discourse context and within-sentence structural and lexical biases, constraints about common ground can be modeled as having weights that depend on their availability and relevance to the task, and are integrated with each other in parallel (e.g., Hanna, Tanenhaus, & Trueswell, 2003; Brown-Schmidt & Hanna, 2011; Brown-Schmidt, 2012). For example, conversational partners can integrate different sources of information from the discourse context, such as the type of discourse prompt and the partner's feedback. Such information may be combined to model common ground in a gradient fashion (see Brown-Schmidt, 2012) rather than all-or-nothing, perhaps even tracking the strength of evidence with which information has been grounded (Clark & Schaefer, 1989). Recent work

(Yoon & Brown-Schmidt, 2019) demonstrated that speakers can combine at least two different kinds of information (prior knowledge and visual perspective) in order to design referring expressions appropriate for two addressees at once during multi-party conversation. However, in that study, speakers did not have to track partner-specific information in memory, as one addressee was always the knowledgeable one, and the speakers could see each addressee's visual perspective displayed on a screen. Further evidence of speakers combining distinct sources of information is that the spatial perspective they adopt is predicted by the convergence of cues, including their own perspective, the partner's spatial perspective, and the intrinsic orientation of the spatial configuration being described (Galati & Avraamides, 2015; Schober, 1993).

A constraint-based model with multiple constraints can be considered an extension of a one-constraint (or "one-bit") model (Brennan, Galati, & Kuhlen, 2010; Galati & Brennan, 2010). In such a model, one or more constraints in memory (most relevant to the task at hand) guide the design or interpretation of utterances according to a conversational partner's knowledge or informational needs. In a similar fashion, single or multiple interacting constraints can be modelled in dynamical models of partner-specific adaptation. For example, Duran and Dale (2014) developed a dynamic "bistable attractor landscape" model of perspective-taking in which a task-relevant constraint (attributions about the partner, such as whether the partner knew the language user's spatial perspective) was represented as a control parameter. The language users' dynamic behavior, in experiments and in simulations, revealed the emergence of bistable states that corresponded to other-centric and egocentric responding, with the control parameter determining the "tilt" of the system that biased behavior toward one perspective or another. In that work, the representation of partner-specific information resulted in a binary distinction in stable patterns of language use. More recently, Duran and Dale (2014)'s dynamical model has

been extended to represent the interaction of multiple constraints operating at different timescales (e.g., “slow” vs. “fast”), supporting predictions about the timing of perspective-taking (Dale et al., 2018).

The proposal that speakers integrate simple constraints, accumulating information at different timescales, contrasts with the view that partner-specific adaptation is driven by the rapid inducement of cognitive alignment across partners through mere exposure or priming. This view, expressed in Pickering and Garrod’s (2004) interactive alignment model, posits that interpersonal alignment across different levels of linguistic representations (from the phonetic and phonological levels, to the lexical and syntactic levels, to their semantic representations and situation models), occurs automatically via priming. In this view, conversational partners need not track anything specific about each other, their interaction, or their common ground, but manage to achieve converging mental representations simply by using their own knowledge and recent experiences as a proxy, under the assumption of parity of production and comprehension. To the extent that processing appears to be partner-specific, this is due to the coincidence of the speakers’ and addressees’ perspectives. In addition, the interactive alignment model focuses only on linguistic information—the linguistic forms previously produced—without specifying how shared visual context could contribute to planning or interact with the effect of shared linguistic forms (Pickering & Garrod, 2004, Figure 2, p. 176; see also Pickering & Garrod, 2013). Therefore, according to the interactive alignment model, when Anna reminisces with Beth about a painting, she should be primed by her more recent description of that painting, regardless of whether and how she viewed or discussed it specifically with Beth or with Camila.

Goals and predictions of current study

In a study that parallels the circumstances of the museum scenario, we manipulated linguistic and visual co-presence independently, to examine whether the modality of co-presence influences not only the grounding of initial referring expressions, but also re-referring later with the same (vs. a different) conversational partner. In our study, speakers arranged geometric shapes that were hard to lexicalize (tangrams cards) with two partners across a total of four rounds. We manipulated whether speakers had previously shared items linguistically and visually, only linguistically, only visually, or not at all during their first interactions with a given partner. Initially, the same items were shared differently with each of the two partners. This experimental design enabled us to address the following questions:

Question 1: How might initial co-presence conditions shape encoding in memory?

To understand impacts of communication modality on partner-specific cues in memory, we began by establishing what occurs during encoding. How do conversational partners coordinate under different co-presence conditions in the first place? Co-presence influences grounding, with shared affordances or co-presence resulting in more efficient coordination (e.g., Brennan, 2005; Clark & Brennan, 1991; Clark & Wilkes-Gibbs, 1986).

We predicted that speakers' descriptions of items that were fully shared (i.e., when both partners had linguistic and visual access) would be more succinct (with fewer idea units, words, and reconceptualizations) than descriptions of items that were shared only linguistically (when only one conversational partner could see the item). We also expected that initial interactions about fully shared items would be more efficient (taking fewer turns and yielding fewer errors) than those shared only linguistically.

Question 2: What are the effects of initial co-presence on partner-specific adaptation?

This is our main research question: When a speaker has previously discussed the same referent with two different partners in two different modalities, do subsequent referring expressions addressed to one partner or the other show any evidence that partner-specific distinctions were encoded? In other words, can the speaker's co-presence conditions with each addressee at initial referring provide constraints strong enough to distinctly shape subsequent referring expressions addressed to each partner?

If distinct associations between perceptual copresence and partner identity are encoded in memory, this should shape how a speaker re-refers to a particular item with two different partners. Specifically, more definite expressions should be used in re-referring when the item had been previously discussed with the partner who had linguistic co-presence than the partner who had no linguistic co-presence. Another well-established process that occurs during repeated referring is attenuation of referring expressions; therefore, there should be fewer idea units (likewise, words, reconceptualizations, and hedges) when an item has been previously discussed (with one partner) vs. not (with the other).

Moreover, if the effects of visual and linguistic co-presence are compounded (e.g., as two interacting constraints), speakers should distinguish information shared both linguistically and visually from that shared through either modality alone. Compounded effects of visual and linguistic co-presence would be consistent with findings that people learn and remember information better when they have experienced it multimodally than when they have experienced it in only one modality (see Shams & Seitz, 2008, for a review). For instance, concrete words paired with relevant visual information are remembered better than other words (Anderson & Bower, 1973; Paivio, 1986), and similarly, stimuli with multisensory pairs (e.g., having an

audiovisual vs. only visual presentation) are categorized more accurately as already-seen (Murray et al., 2004).

In contrast to these predictions, the interactive alignment model predicts that only the speaker's most recent linguistic experience (what they described most recently), regardless of shared visual copresence, should prime subsequent referring. That is, in this view, speakers would mark items as given or known based their own experience: there would be no distinct partner effects in re-referring, except that the experience with the most *recent* partner should be more activated and more likely to shape subsequent interaction.

Question 3: What are the effects of initial co-presence on subsequent coordination and task success

Does the modality of previously shared information influence how conversational partners coordinate referring during subsequent conversation? To address this question, we examined whether efficiency and accuracy were impacted by the initial modality of co-presence with that partner.

Our approach

To investigate Questions 1 and 2, we used measures that captured the content of a referring expression, definiteness (whether an expression was marked as old), its provisionality (whether an expression was tentatively presented as new), and its ease of planning. For content, we examined the number of idea units and words speakers produced. Speakers have been shown to use more words (Galati & Brennan, 2010) and idea units (Horton & Gerrig, 2002) to describe information that is new to a conversational partner than information that is old. We also assessed the number of perspective-switches (reconceptualizations), indicated by the number of new idea units. Speakers have been shown to use more reconceptualizations with a new partner than with a

previous partner when they can keep track of what categories of information they've previously talked about (Horton & Gerrig, 2005b, 2002). Here, we assess not only the effect of linguistic co-presence, but also of visual co-presence and its interaction with linguistic co-presence on these key aspects of propositional content. For definiteness, we counted definite expressions and meta-comments, as speakers have been shown to be more likely to mark explicitly a referent as having been shared with a partner (e.g. "*from the last time*") the more frequently they have referred to it with their partner (Brennan & Clark, 1996). For the provisionality of speakers' descriptions, we counted hedges (expressions such as "kind of" and "almost"). Such hedges could be sensitive to perceptual co-presence; the less certain speakers are that their partners will accept how they conceptualize a referent, the more likely they are to use hedges (e.g., Brennan & Clark, 1996; Horton & Gerrig, 2002). Finally, for ease of planning, we examined speakers' description onset latencies, which are thought to reflect the ease with which speakers establish whether an item was previously shared with a particular partner (e.g., Horton & Gerrig, 2005b; Hwang, Brennan, & Huffman, 2015).

To assess Question 3 (communicative efficiency), we examined the number of conversational turns and the pairs' errors across the entire dialogue, following evidence that length of the dialogue and task accuracy are sensitive to the affordances of the interaction (Brennan, 2005; Clark & Krych, 2004; Gergle, Kraut, & Fussell, 2004).

Method

Design and task

In a referential communication task, a director matched sets of picture cards (with a total of eight cards per set) in repeated rounds with one of two matchers (Matcher A and Matcher B). The rounds were divided into two phases: Phase 1 was used to allow the director to establish

common ground under different conditions of co-presence with each of the two matchers in two separate rounds, and Phase 2 was used to observe, in two final rounds, whether and how the director adapted the referring expressions addressed to each matcher. In other words, the director of each triad shared the same cards with both matchers, but in different co-presence conditions; therefore any partner-specific adaptation in Phase 2 would depend on the director's ability to distinguish how the same card had been encoded with each matcher in Phase 1.

Phase 1

Typically in a matching task, conversational partners have identical sets of cards that they must arrange in a target order, while visually separated by a barrier (e.g., as in Clark & Wilkes-Gibbs, 1986). To set up the co-presence conditions in Phase 1, we departed from this convention in two critical ways: (1) we occluded only the partners' private staging areas while having them place cards in a target area visible to both, and (2) we had the director match a systematically different subset of each card set with each matcher in Phase 1.

This subset was distributed as follows: Two cards were present in both the director's and the matcher's sets (so were shared both visually and linguistically), two were in only the director's set (to be shared only linguistically, with the director describing them to the matcher), two were in only the matcher's set (to be shared only visually, when the matcher placed them in the target area), and the two remaining cards were missing during Phase 1 (Table 1 illustrates this setup). Critically, the co-presence condition for each card in a given set varied across matchers: Cards that were shared linguistically *and* visually with Matcher A were completely absent for Matcher B, and vice-versa; and cards that were shared only linguistically with Matcher A were shared only visually with Matcher B, and vice-versa. As a consequence of this feature of the design, the director described each card only once in Phase 1, to either Matcher A

or B. That is, by the end of Phase 1, all items were equally linguistically “given” for the director, since each card had been described once. In a given set, 4 of the 8 cards were described to Matcher A (two shared fully and two shared only linguistically) and 4 cards were described to Matcher B (again, two shared fully and two only linguistically), with the co-presence status of each being complementary across the two matchers, as shown in Table 1.

To establish the linguistic-only and visual-only co-presence conditions, respectively, participants followed two rules in Phase 1: (1) When the director described a card that the matcher determined was lacking, the matcher placed a “joker” (placeholder) card in the target location, and (2) when, at the end of the round, the matcher was left with the two cards that director lacked, the matcher silently handed these cards to the director, who silently confirmed them as missing and placed them in the last two available slots of the target area (slots 5 and 6). By the end of a round in Phase 1, the 6-slot target area was filled with two cards that had been shared linguistically and visually, two joker cards standing in for items shared only linguistically, and in the final two slots, two cards shared only visually (and not discussed).

Phase 2

In Phase 2, the director matched all eight target cards with each matcher privately (Rounds 3 and 4). As in Phase 1, partners could interact freely and correct any errors. The order in which matchers returned for Phase 2 was counterbalanced: Either Matcher A returned first for Round 3 in Phase 2, followed by Matcher B (matcher order ABAB), or else Matcher B remained for Round 3 in Phase 2, with Matcher A returning for Round 4 (order ABBA). Because B was always the most recent partner from Phase 1, finding stronger partner-specific effects for Matcher B than Matcher A with respect to linguistic co-presence would suggest that the effect is driven primarily by priming, as predicted by interactive alignment.

Counterbalancing

Each triad of participants followed this two-phase design twice, with two different sets of cards displaying tangrams. Each triad matched one set of tangrams objects in Matcher order ABAB and the other in order ABBA, counterbalancing their pairing and order across triads². Co-presence varied both within-triads and within-items; the tangram cards were rotated across four lists so that each individual card appeared equally often in all four co-presence conditions in Rounds 1 and 2 (Phase 1) and in each of the eight target slots in Rounds 3 and 4 (Phase 2). Each triad experienced one list, and the co-presence condition of each tangram in a given list was complementary across the two matchers (see Table 1). These pairings of Phase 1 co-presence conditions were rotated across the four lists: in one list a given item was shared fully (LV) with Matcher A and not with Matcher B (so was new (N) in Phase 2); in another list that card was N with Matcher A and LV with Matcher B; in a third list the card was shared only linguistically (L-only) with Matcher A and only visually (V-only) with Matcher B, and in a final list it was V-only with Matcher A and L-only with Matcher B. The four lists of items per tangram set, along with the counterbalancing of the items across positions, can be found in the Materials folder of our OSF repository for the project (<https://osf.io/7dhk5/>).

Participants

Ninety-six students from Stony Brook University participated in 32 triads: 32 participants served as Director and 64 as Matchers. All identified themselves as native speakers of English; 58 were female and 38 were male. Participants within a triad were assigned to their roles

² Each triad also completed this two-phase design twice more, with two additional sets of cards displaying common objects that are easy to label (a total of 16 line drawings from different categories). These were used to address a separate question regarding whether the modality of co-presence would influence the articulation of referring expressions. The order of the type of cards used (tangrams and common objects) was counterbalanced across triads.

randomly. They received either research credit that they could use to satisfy a requirement for a psychology course or prorated payment at the rate of \$9/hour.

[INSERT TABLE 1 ABOUT HERE]

Table 1

Distribution of items (cards) in Phase 1 and the resulting co-presence conditions for sharing with either partner. Jokers were placeholder cards.

Distribution of items in Rounds 1 and 2 (Phase 1)		
	Round 1 with Matcher A	Round 2 with Matcher B
Director has	card 1, card 2, card 3, card 4	card 5, card 6, card 7, card 8
Matcher has	card 1, card 2, card 5, card 6, 2 jokers	card 3, card 4, card 7, card 8, 2 jokers

Resulting information status of items in Rounds 3 and 4 (Phase 2)		
	with Matcher A	with Matcher B
card 1 & 2	shared linguistically and visually	not shared (new)
card 3 & 4	shared linguistically only	shared visually only
card 5 & 6	shared visually only	shared linguistically only
card 7 & 8	not shared (new)	shared linguistically and visually

Note. Card numbers refer to the identity of cards, not to the order in which they were matched.

Procedure

Participants were told that the study investigated collaboration, and drew cards to determine their roles as Director, Matcher A, or Matcher B. The rules of the matching task were then explained.

During each round, the director-matcher pair sat at a table facing one another while separated by a barrier that occluded the matcher's staging area and the director's stack of cards, but not each other's faces (Figure 1). Both partners could see the target area and reach its slots. A digital camcorder on a tripod behind the director captured the director's cards, the director's

view of the matcher, and an oblique view of the target area, while another digital camcorder provided a direct view of the target area and side views of the director and matcher. A high-quality audio recorder recorded speech from the director's headset microphone.

[INSERT FIGURE 1 ABOUT HERE]

Figure 1

Arrangement of a director and a matcher in a Phase 1 round. Phase 2 rounds involved the same arrangement, except the target area was a board with eight slots. (Inset is the view from the camera over the director's shoulder; the 2 cards with reddish figures are jokers.)

At the start of each round, the director's cards were stacked in a prearranged order, whereas the matcher's cards were spread out in front of them in their private staging area. Participants were told that the goal of the task was for the matcher's cards to be placed on the target area (a board with numbered slots), matching the target order of the director's stack of cards. Directors were instructed not to turn over the next card in their stack until the matcher had placed the current card on the target area. Participants were told that they could correct any errors they made. For instance, if the matcher placed the wrong card on the board, the director could correct them. Or if the matcher could not identify the card described by the director but was out of joker cards, they could backtrack to discuss whether a joker card had been mistakenly used earlier.

The structure of the matching rounds was described: participants were told that, for each set of cards, they would be playing for four rounds divided into two phases. In the first two rounds (Phase 1), the director would be matching a subset of the cards for one round with Matcher A and one round with Matcher B, separately, and that special rules for those two rounds applied, since their subsets of cards would not be identical. The two special rules of the first two rounds (Phase 1) were explained. Participants were told that in the last two rounds of playing with a given set of cards (Rounds 3 and 4), they would be matching cards from the previous rounds; for Rounds 3 and 4, the director and each matcher would have the same set of cards, and the special rules of the first two rounds no longer applied. All participants were asked to ignore the experimenter, who remained in the room in order to operate the recording equipment, give pairs their cards for each round, and change the target area boards between Phases 1 and 2.

After these instructions, the director completed two Phase 1-style practice rounds with tangram cards, one with each of matcher. Next, the triads carried out all experimental rounds of

the matching task, while being video and audio-recorded. While the director matched cards with one matcher, the other matcher waited outside the room. After all the matching rounds were completed, the director and matchers were administered a questionnaire³ about the items. Finally, they were all debriefed.

Materials

Two copies apiece of 16 tangrams (pictures of geometric objects that are not associated with lexicalized labels) were printed on 3 x 5 inch index cards. The 16 tangrams were divided into two sets of eight cards: one set resembling human figures, the other resembling animals (see our online repository for the full sets: <https://osf.io/7mxy8/>). Each of these sets was used across four rounds. For each of the two sets, an additional tangram (a human-like figure and an animal-like figure) served as the matchers' ninth card in Phase 2. This forced pairs to discuss all the critical items and not refer to any critical item as "the last one". The tangrams were selected to have comparable naming times according to a norming study by Swets, Jacovina, and Gerrig (2013). In addition, two images of jokers were printed on 3 x 5 inch index cards to be used by matchers in Phase 1.

³ The questionnaire asked participants to recall how each item had been shared with their partners during Phase 1. For each tangram, directors had to choose one of the following options with respect to Matcher A: (a) *both I and Matcher A had it*, (b) *only I had it (Matcher A used a joker)*, (c) *only Matcher A had it (A placed it in slot 5 or 6)*, (d) *neither I nor Matcher A had it*, (e) *I don't remember*. The same options with respect to Matcher B were listed for each item as well. Matchers were also administered adapted versions of this questionnaire, but only the directors' responses were examined, as we were interested in their explicit memory for how they had shared the items with each partner. We do not report the questionnaire results here because the responses were subject to interference from Phase 2. Since questionnaires were administered at the end of the experiment, directors had already encountered every item linguistically and visually with both matchers. We did not administer questionnaires at the end of Phase 1, as we didn't want directors to be explicitly queried about the status of information of each item prior to interacting with matchers in Phase 2. Not surprisingly, given the time point at which the questionnaires were administered, directors showed a bias in reporting that a tangram had been shared both linguistically and physically in Phase 1. Although they correctly identified 86% of the tangrams that had been shared linguistically and physically with a given partner, they also incorrectly identified 56% of the remaining tangrams as having been shared linguistically and physically with a given partner. This does not preclude the possibility that during Phase 2, directors assessed accurately how they had shared an item with a given partner in Phase 1.

The target area for Phase 1 was a 12 x 36 inch board with six slots, numbered 1 through 6. Slots 5 and 6 of this board, where directors placed the visually-only shared cards towards the end of a round in Phase 1, were highlighted in yellow. The target area for Phase 2 was a 12 x 40 inch board with 8 slots, numbered 1 through 8.

Transcription

For each triad, all four rounds of each set of cards were transcribed in detail, including spoken contributions by both matchers and directors. The transcripts included fillers such as *uh* or *um*, pauses, interruptions (both directors' self-interruptions and interruptions by matchers) and restarts. Instrumental actions (e.g., the matcher placing a card on the board) and non-verbal feedback such as head nods and facial displays (e.g., frowning in confusion) were also annotated.

Coding

For Phase 1, we sought to establish whether directors and matchers coordinated differently over cards that were linguistically and visually shared vs. those that were shared only linguistically (Question 1). Toward that end, we coded the *total* amount of content (number of idea units, words, and reconceptualizations), the incidence of provisionality (hedges) across the entire description produced by the director, and the number of turns and errors within the interaction.

For Phase 2, we sought to examine whether directors adapted *initial descriptions* addressed to a partner according to how they had shared information previously. Toward that end, in the director's initial description in Phase 2, we examined the amount and type of content (number of idea units, words, and perspective-switches as captured by the number of reconceptualizations), provisionality (the presence of hedges), and definiteness (the presence of

definite expressions and meta-references). Initial descriptions were defined as the director's first turn, before any overt contribution in the form of explicit feedback from the matcher⁴.

By focusing on the director's initial descriptions, we aimed to tap into partner-specific adaptation that was *not* performed as a repair in response to the matcher's feedback, but as part of initial speech planning. This does not preclude the possibility that, during their initial description, directors still monitored their matcher's non-verbal behavior by looking at the matcher's face over the barrier and adapted their linguistic behavior accordingly. To address this possibility, we also measured the time directors took to produce the first content word of a Phase 2 description (description onset latency), which we took to reflect adaptation that was even less affected by any feedback.

Finally, we examined whether pairs in Phase 2 coordinated differently according to how information had been shared in Phase 1 (Question 3). Taking as a unit of analysis the entire exchange between the director and matcher for a given item, we considered the number of conversational turns for the entire item description in Phase 2 and the number of errors for a description in Phase 2. We also coded the type of the matcher's contribution in response to the director's initial description in Phase 2; as these measures were not of primary interest, we present them as an online supplement in our repository (<https://osf.io/wfbjr/>).

Below, we describe in more detail each of our dependent measures.

Content: Number of words

⁴ There were two exceptions to this definition of the directors' initial descriptions, relating to cases where the matcher's speech overlapped the director's. One exception was when the matcher attempted to speak during the director's first turn, but the director continued speaking without pausing or interrupting herself. Typically, matchers stopped speaking almost immediately in these cases. The other exception was when matchers provided feedback (e.g., saying "mm-hmm" or laughing) while the director spoke without pausing or interrupting herself in response to this feedback. In these cases, we took the directors' continued speaking to indicate that they had not yet finished with their initial description, and considered matcher's next contribution to mark the end of the director's initial description.

We counted the number of words in the directors' initial tangram description in Phase 2, as well as their total number of words for each tangram in Phase 1. Contracted words (e.g., *not*, *is*, *have*), audible feedback responses (*mm-hmm*, *nuh-uh*) and expressive exclamations (e.g., *wow*, *oh*, *ugh*) were counted as separate words. Words interrupted mid-word were excluded from the word count.

Content: Number of idea units

Because word counts may be inflated by hedging (see below), we also considered the number of idea units in speakers' initial descriptions. This was our principle measure for assessing content. We coded each content word (i.e., each noun, adjective, action verb, adverb, but also prepositional phrases) as belonging to a separate idea unit (similar to Horton & Gerrig, 2002). Verb phrases encoding directionality (e.g., "*facing to the left*", "*pointing to the left*") were counted as single idea units just like prepositional phrases (e.g., "*to the left*"). Compound nouns (e.g., "*stop sign*", "*lava lamp*") were also counted as single idea units.

Content: Number of reconceptualizations

We defined reconceptualizations as those idea units in the director's initial description to a matcher in Phase 2 that did not overlap with those of their description in Phase 1 (i.e., were unique to Phase 2). Identical idea units were classified as overlapping, but synonyms or expressions that were semantically equivalent were also coded as overlapping (e.g., *large-big*, *somebody-person*, *sideways-to the side*, *facing to the left-looking to the left*).

For each description of an item in Phase 2, reconceptualizations were coded relative to that item's description in Phase 1. Note that each item was described only once in Phase 1, but twice in Phase 2 (once to each partner). Appendix A includes an example of coding idea units and reconceptualizations for a card that a director shared linguistically with Matcher A and

visually with Matcher B in Phase 1. That example illustrates how reconceptualizations in initial descriptions to Matcher A and to Matcher B in Phase 2 are coded relative to the description of that item (to Matcher A) in Phase 1.

Provisionality: Presence of hedges

The following kinds of expressions, marking descriptions as provisional, were counted as hedges: kind of, sort of, maybe, probably, possibly, almost, in a way, in some sense, a little bit, if anything, -looking thing (e.g., “*shield-looking thing*”), -type (of) thing (e.g., “*foot-type thing*”), -ish (e.g., “*triangle-ish*”, “*back neck-ish*”), -y (e.g., “*rhombus-y*”).

Definiteness: Presence of definite expressions and meta-references

We coded for definiteness by identifying definite expressions and meta-references indicating prior linguistic or visual co-presence. Definiteness included the presence of definite expressions (e.g. “*the person falling*”, “*the trumpet player*”), noun phrases preceded by definite pronouns (e.g., “*our clown*”, “*our morphed creature*”), and meta-references (e.g., “*it’s already come up before*”, “*that duck character we talked about earlier*”). Definite expressions arising from shared community membership, based on historical or other shared socio-cultural knowledge, were not counted as definite expressions (e.g., “*the Batman symbol*”, “*the presidential seal*”, “*the geese that you see on campus*”, or “*those monkey people with those pointy hats in the Wizard of Oz*” as introductory mentions). This was because we wished to isolate the effect of shared experience due to co-presence, over and above community membership⁵.

⁵ Speakers’ references could, of course, reflect both community membership and co-presence (linguistic or visual). There were some such instances in our corpus. For example, a director could say “this looks like the Batman symbol” in Phase 1 and “this is our Batman symbol” with the same matcher in Phase 2. Since we wished to isolate the effect of shared experience due to co-presence over and above community membership, the first instance (“the Batman symbol”) would not count as a definite expression, but the second instance (“our Batman symbol”) would.

Planning: Description onset time

To measure how long it took directors to initiate their descriptions, we used the video annotating tool ELAN (Brugman & Russel, 2004) to identify two time points in the digital video-recording for descriptions in Phase 2: (1) the first video frame in which the tangram on the current card became fully visible to the camera behind the director (note that directors were instructed not to turn a card until the matcher had placed the previous card in the target area), and (2) the onset of the first content word that was part of the description of the current card.

Task efficiency: Number of turns

Uninterrupted stretches of speech by a director or matcher were counted as turns. Instrumental actions performed by participants, such as matchers visibly placing cards on the target area were also counted as turns, even when these actions were performed in silence. If a matcher ratified their acceptance of a director's description in speech (e.g., by saying "*got it*", "*yeah*", or "*mm-hmm*") while simultaneously placing their card selection on the target area, their speech and instrumental action were coded as belonging to the same turn. Exchanges at the beginning or end of a round that dealt with readiness to begin (e.g., D: *ready?* M: *yeah*) or concluding the round (e.g.: D: *and that's it*) were excluded so as not to inflate the turn count of the first and last items. When non-verbal feedback, such as head nods or facial displays, were unaccompanied by speech, they were counted as a turn only if the conversational partner was attending to that feedback⁶.

⁶ The rationale of this decision was to code and include as turns only those nonverbal signals that had the potential to be communicative in the absence of accompanying speech. In this task, partners did not continually monitor each other: partners often attended their individual task space, which was occluded from the other person's view (e.g., the matcher looking at the cards in front of them; the director focusing on the card they were describing). This meant that the nonverbal signals of one partner could be missed by the other in some cases. In those cases, we reasoned that the unattended contribution (e.g., an unattended head nod) did not advance the dialogue and should not be counted as a conversational turn. Following Brennan and Williams (1995), we focused on signals that had to be not only informative (e.g., indicating acceptance or agreement), but also could be apprehended and processed by the addressee. Non-verbal feedback was coded by referring to the two videos available from each session (as shown in

Task accuracy: Number of errors

We coded for three possible types of errors in matching tangrams, the last two being specific to Phase 1: (1) selecting the wrong tangram, (2) selecting a tangram instead of a joker card, and (3) selecting a joker card instead of a tangram.

Reliability

Reliability was established for those measures involving judgments about the equivalency of units of meaning (i.e., idea units and reconceptualizations) and about segmentation (i.e., description onset times). For idea units and reconceptualizations, a second coder (a research assistant unaware of the matchers' knowledge status and of the hypotheses) redundantly coded four randomly chosen triads, corresponding to approximately 12.5% of the corpus. This resulted in a total of 192 judgments for idea units (16 items per triad, each item described 3 times) and 128 judgments for reconceptualizations (16 items per triad, with each item described twice in Phase 2). To ensure coding was blind to how information had been previously shared between participants, the matchers' identity in Phase 2 was disguised in the transcripts. The coders' (the first author and the second coder) judgments exhibited strong agreement: They identified exactly the same number of idea units for 77% of the cases and differed by only one idea unit for another 16% of the cases. The pair of coders identified the same number of reconceptualizations for 74% of the cases and differed by only one reconceptualization for another 21% of the cases.

Krippendorff's alpha (Hayes & Krippendorff, 2007) was high for both idea units (.98) and reconceptualizations (.86). The first author's coding was used for subsequent analyses.

Figure 1), relying primarily on the view of the camera stationed behind the director (i.e., gauging whether the matcher's feedback could be perceived by the director).

Description onset latencies were coded by four coders (the first author and three undergraduate research assistants). To assess reliability, all coders redundantly coded approximately 9% of the corpus (three triads), resulting in a total of 91 measurements of description onset. The mean absolute differences between the first author's and each of the research assistants' measurements were small (37.25 ms, 4.32 ms, and 87.09 ms)⁷. The differences arose primarily from disagreement about when the image on the director's rotating card first became fully visible (time point 1) as opposed to disagreement about the beginning of the first content word (time point 2). Overall, there was high agreement among coders on the measurement for onset latencies (Krippendorff's $\alpha = .94$).

Statistical Models

Data were analyzed with generalized linear mixed-effects models, which included random effects for participants (directors or triads, as relevant) and items (tangram cards) in each model. The models were fitted using the lme4 package (Bates, Mächler, Bolker, & Walker, 2015) in R (R Core Team, 2016). We built maximal models with a full random effect structure (as long as the model converged), following the recommendations of Barr, Levy, Scheepers, and Tily (2013). If a model failed to converge, we simplified it by removing terms from the random effect structure in a theoretical motivated manner starting with the higher order terms of the least theoretical interest⁸ until the most complex model that converged was obtained (following the

⁷ For the 3 triads for which onset times were coded redundantly for reliability, the first author's coding was used. For the remaining triads, the individual coder's measurements were used. These included the coding of 4 additional triads by the first author, 8 triads by coder 2, 10 by coder 3, and 5 by coder 4. (As we describe in the Statistical Models section, two triads were excluded from coding due to missing or uncodable data; see footnote 9.)

⁸ When models failed to converge, we sequentially removed the random slopes for: (i) the interaction between partner order and partner identity (ii), the interaction between partner identity and visual co-presence, (iii) the interaction between partner identity and linguistic co-presence, (iv) partner order, (v) partner identity, (vi) visual co-presence, and finally (vii) linguistic co-presence.

recommendations of Bates, Kliegl, Vasishth, & Baayen, 2015). The results from the first converging model are reported.

For each dependent variable for Phase 2 models⁹, we entered as fixed effects, the relevant types of co-presence, partner identity (Matcher A vs. Matcher B), partner order (ABAB vs. ABBA), along with those interactions that were of theoretical interest. For Phase 2 models, co-presence type was coded as two binary variables: linguistic co-presence (item previously shared through linguistic co-presence vs. not) and visual co-presence (item previously shared through visual co-presence vs. not). We included four interaction terms in the models of Phase 2: the two-way interaction between linguistic and visual co-presence, the two-way interaction between partner identity and partner order, and the two-way interaction of each co-presence constraint with partner identity. The interaction between linguistic and visual co-presence addresses our key theoretical question about whether the effect of one constraint is amplified by the other during re-referring. The interaction between partner identity and partner order captures any order effects of practice or fatigue: it permits evaluating, for example, whether there is more attenuation in references addressed to partner A in order ABBA than ABAB.

The interactions between co-presence constraints and partner identity capture whether adaptation in Phase 2 is more pronounced for the more recent partner (Matcher B). Of these two potential interactions, the one between linguistic co-presence and partner identity is of theoretical

⁹ The Phase 2 dataset included 1020 data points (LV: 225, L-only: 255, V-only: 256; New: 254), distributed across 32 triads and 16 items, following the exclusion of four item descriptions that were missing or uncodable. The dataset of description onset times had a different structure because it was derived from video recordings rather than transcripts. It excluded data from two triads: for one triad the video-recording of the camera behind the director had been corrupted; for the other triad, the director, despite instructions, previewed all cards multiple times before describing them. In addition, 95 individual trials were excluded because directors either turned over their card before the matcher placed the previous card on the target area or continued to discuss the previous item while viewing their current card. Thus, the description onset times dataset included a total of 865 observations, across 30 triads and 16 items.

interest; since B is always the most recent partner when Phase 2 begins and A never is. The interactive alignment model would predict an interaction between linguistic co-presence and partner identity (with more attenuation for B than A), suggesting that priming rather than partner-specific encoding is what drives any attenuation effects.

Given the number of parameters of theoretical interest in the models of Phase 2, we applied a Bonferroni correction to control the familywise error rate. The 4 most critical effects of theoretical interest (the fixed effects of linguistic co-presence, visual co-presence, their interaction, and the interaction between linguistic copresence and partner identity) constitute the key “family” of tests, as they permit us to distinguish between theoretical accounts (see Lakens, 2016; Gelman et al., 2012). This conservative Bonferroni correction resulted in an adjusted alpha of .0125.

For Phase 1 models¹⁰, co-presence was coded as having two levels: items shared linguistically and visually (LV) versus items shared linguistically-only (L-only); this distinction captures the binary visual co-presence constraint, while keeping the linguistic co-presence constraint constant. For Phase 1 models, partner order wasn’t included in the models since the two rounds of Phase 1 always had the order, AB. For Phase 1 models, we applied a conservative Bonferroni correction for the 3 model parameters: the effect of co-presence (LV vs. L-only), partner identity (A vs. B), and its interaction with co-presence. The effect of co-presence was of primary theoretical interest, but the effect of partner identity and its interaction with partner order could reveal theoretically relevant practice effects (e.g., more efficient interactions and more concise descriptions with matcher B). This resulted in a Bonferroni adjusted alpha of .0167.

¹⁰ The Phase 1 dataset included 511 data points for each dependent measure (LV: 256; L-only: 255), distributed across 32 triads and 16 items, following the exclusion of one item with an uncodable description (the director identified this item by mentioning that the matcher had previously selected it in error).

As random effects, we had intercepts for directors (or triads, as relevant) and items, as well as random slopes, to account for between-participant and between-item variation; for the effect of co-presence constraints, partner identity, partner order (as appropriate); and for the interactions, which served as fixed effects, for directors (or triads) and items, as long as the model converged. Appendices B and C (for Phase 1 and Phase 2, respectively) include tables reporting the variance and standard deviations for each random effect in each model. The random effect structure of each model can be inferred from the random effects reported.

The *p*-values were obtained from the lmerTest package (Kuznetsova, Brockhoff, & Christensen, 2016) using the Satterthwaite's method. Captured variance of overall models is reported as Conditional R^2 variance explained by fixed and random factors together, which was computed using the MuMIn R statistical package (Johnson, 2014). For binary dependent variables (the presence of definite expressions and hedges), we used mixed logistic regression models with binomial error structure.

When visual inspection of residual plots for the models revealed deviation from normality and homoscedasticity, the dependent variable was log-transformed. That was the case for total number of words (Phase 1 and 2), total number of idea units (Phase 1 and 2), and number of turns (in Phase 2).

Data, Materials, and Code Sharing

De-identified transcripts, raw data files (of coded transcripts and ELAN files for coded description onsets), along with code for preparing the data, obtaining descriptives and visualizations, specifying planned contrasts, and testing the statistical models, are available through our OSF repository for the project (<https://osf.io/7dhk5/>). The raw video and audio data are not included in this repository due to privacy restrictions. However, videos of pairs for which

both participants signed release forms can be shared upon request to the first author. Card stimuli, counterbalancing lists, design protocols, and supplementary analyses (e.g., on the matchers' feedback) and visualizations are also included in the repository.

Results

Effects of co-presence on grounding during initial interaction (Phase 1)

We first examined whether directors in Phase 1 described items that were fully shared differently from those that had been shared only linguistically. We expected that pairs would have more difficulty grounding items that were not fully shared, and that this would be reflected in the descriptions. As predicted (and consistent with previous studies), when directors described items that their matchers didn't have, their descriptions had more words and idea units, were more likely to contain hedges, and unfolded over more turns compared to when matchers had the item being described (see Table 2). Pairs also made four times as many errors when discussing items that the matchers didn't have. Out of the total of 62 errors, matchers selected a tangram card instead of the joker card 77.4% of the time, the wrong tangram 14.5% of the time, and a joker card instead of a tangram they actually had 8.1% of the time.

[INSERT TABLE 2 ABOUT HERE]

Table 2

Means and standard deviations—across the two conditions of co-presence in Phase 1—sharing linguistically and visually (LV) or only linguistically (L-only)—for directors' total number of words, total number of idea units, and proportion of items with hedges; for the pairs' total number of turns and number of errors (per item).

	LV		L-only	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Number of words	36.20	29.63	62.19	46.42
Number of idea units	8.85	6.46	15.13	11.46
Proportion of items with hedges	0.46	0.50	0.59	0.49
Number of turns	4.86	3.85	7.31	6.07
Number of errors	0.05	0.28	0.20	0.48

The linear mixed effects models confirmed significant differences in language use between the two co-presence conditions (see Appendix B). Co-presence condition (sharing items linguistically and visually vs. only linguistically) was a significant predictor of the director's language use in terms of the total numbers of words, the number of idea units, and the incidence of hedges. Co-presence was also a significant predictor of the pair's performance, predicting the total number of turns to coordinate over matching an item and the number of errors made while doing so.

The interaction between co-presence condition and whether the partner was first or second (partner identity A or B) also significantly predicted total words and idea units. Whereas directors used comparable numbers of words and idea units with both partners when items were fully shared (the differences between Matcher A and B being only 1.77 for words and 0.89 for idea units), they used more words and idea units with matcher A (i.e., in the first round) when items shared only linguistically (the differences between A and B being 18.70 words and 4.47 idea units). It may be that as directors became more familiar with the challenge of describing referents that were hard to lexicalize, they provided more effective descriptions in the second round (with Matcher B).

In sum, pairs coordinated differently over each item during their first interaction, depending on how they shared the item. Having visual co-presence while discussing these items

shaped the pairs' interactions: when directors described items that their matcher didn't have, they produced more elaborate descriptions (with more words and idea units), and more provisional descriptions (with more hedges) compared to when describing items that the matcher had. Pairs discussed items shared only through linguistic co-presence over more turns and made more errors. These findings are consistent with previous work showing that, when relying only on spoken evidence, conversational partners go through a lengthier process of grounding, or checking that something has been mutually understood (Brennan, 2005; Clark & Krych, 2004; Gergle, Kraut, & Fussell, 2004).

Effects of co-presence on partner-specific adaptation in subsequent interactions (Phase 2)

Having established that partners interacted differently over items that were fully shared compared to those shared only linguistically in Phase 1, we turn to the key theoretical question of whether directors would adapt subsequent referring expressions in a partner-specific way that took the type of previous co-presence with each matcher into account in Phase 2. By representing linguistic and visual co-presence as binary constraints, we assessed whether each co-presence constraint would affect descriptions, and especially whether the two constraints would interact. An interaction would constitute evidence that both sources of perceptual information are retained in memory traces of shared experiences, with their effects being compounded.

Content: idea units, words, and reconceptualizations

First, we examined adaptation in terms of content: the number of idea units, the number of words, and the amount of new content (or reconceptualizations, which suggest a shift in perspective). As shown in Table 3, directors produced numerically fewer idea units, words, and reconceptualizations in their initial description when items had been shared fully than when they had been shared only linguistically. In turn, they produced fewer idea units, words, and

reconceptualizations when items had been shared only linguistically than when had been previously shared only visually. This gradient change in idea units across co-presence conditions is illustrated in Figures 2, as well. The same pattern is observed for words and reconceptualizations (see online supplement at: <https://osf.io/ekxu7/>). Somewhat surprisingly, directors used numerically more words, idea units, and reconceptualizations when referring to items that had been shared visually-only than to items that were brand new to matchers.

As shown in the linear mixed effects models reported in Table 4, directors produced fewer words, idea units, and reconceptualizations when items were previously discussed: the binary constraint for linguistic co-presence was a significant predictor in these models. In contrast, visual co-presence on its own did not significantly predict these measures of content. However, for the models for words and idea units, there was a significant interaction between linguistic and visual co-presence, suggesting that the effect of linguistic co-presence depended on visual co-presence. This interaction is reflected in the pattern illustrated in Figure 2, in that directors attenuated fully shared items more than items shared linguistically-only. For reconceptualizations, the interaction between linguistic and visual co-presence was marginal ($p = .0149$) under our conservative alpha level ($p = .0125$).

Linguistic co-presence did not interact with partner identity: directors did not attenuate their descriptions any more with the more recent partner (B) than with the other partner (A). This is consistent with linguistic attenuation being driven by partner identity rather than by priming and recency (contra interactive alignment). Recency did matter, however, when it came to visual experience; the interaction between visual co-presence and partner identity was a significant predictor in the models for idea units and reconceptualizations. This is in part because when items were shared only visually with Matcher B (the most recent partner), directors produced

more idea units and reconceptualizations (for idea units: $M = 8.74$, $SD = 6.15$ with Matcher B, vs. $M = 6.96$, $SD = 4.79$ with Matcher A; for reconceptualizations: $M = 3.96$, $SD = 4.47$ for B, vs. $M = 2.87$, $SD = 3.43$ for A). Although the recency of experiences from Phase 1 influenced descriptions of visually-only items, this finding does not bear on interactive alignment, since that account doesn't make predictions about the role of visual co-presence.

[INSERT FIGURE 2 ABOUT HERE]

[INSERT TABLE 3 ABOUT HERE]

[INSERT TABLE 4 ABOUT HERE]

[INSERT TABLE 5 ABOUT HERE]

Table 3

Means and standard deviations—across the four conditions of co-presence in Phase 2: items that had been shared linguistically and visually (LV), only linguistically (L-only), only visually (V-only), or not at all (New) in Phase 1—for directors' number of words, number of idea units, and reconceptualizations, the proportion of items with hedges in the directors' initial descriptions, the proportion of items with definite expressions or metareferences, the directors' description onset times (in ms), the pairs' total number of turns and number of errors.

	LV		L-only		V-only		New	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Number of words	17.51	11.06	22.59	14.46	28.27	20.37	23.19	16.17
Number of idea units	5.17	3.22	6.54	4.03	7.85	5.58	6.80	4.76
Number of reconceptualizations	2.02	2.23	2.45	2.65	3.41	4.01	3.03	3.41
Proportion of items with hedges	0.12	0.32	0.18	0.38	0.29	0.45	0.24	0.43
Proportion of items with definite expressions or meta references	0.36	0.48	0.40	0.49	0.21	0.41	0.16	0.37
Description onset times	2993.90	1491.95	3365.65	1749.22	3467.11	1940.19	3244.89	1681.84
Number of turns	2.99	1.87	4.38	3.90	4.69	5.31	4.25	3.32
Number of errors	0.01	0.11	0.05	0.26	0.04	0.24	0.03	0.16

Figure 2

Violin plots representing the distributions of the number of idea units in the directors' initial description in Phase 2, according to the co-presence condition of that item with that partner in Phase 1: shared linguistically and visually (LV), linguistically only (L), visually only (V), or being new (N). Boxplots represent the median and quartiles (Q1, Q3), and the red dot represents the mean of each condition. Black dots indicate observations with values greater than Q3 plus 1.5 times the interquartile range.

Table 4

Mixed-effects models for the total number of idea units (log-transformed), words (log-transformed), and reconceptualizations in directors' initial descriptions in Phase 2.

Predictor	Initial idea units				Initial words				Reconceptualizations			
	R ² = .39				R ² = .44				R ² = .34			
	B	SE	t	p	B	SE	t	p	B	SE	t	p
Intercept	1.68	0.08	22.06	<0.001	2.92	0.08	37.61	<0.001	2.73	0.33	8.16	<0.001
Linguistic Co-presence (Yes vs. No)	0.20	0.05	4.32	<0.001	0.23	0.05	4.90	<0.001	0.99	0.29	3.38	<0.001
Visual Co-presence (Yes vs. No)	0.06	0.03	1.82	0.070	0.04	0.03	1.22	0.221	0.02	0.20	0.10	0.919
Partner (A vs. B)	-0.04	0.03	-1.32	0.187	-0.03	0.03	-0.93	0.354	-0.10	0.16	-0.59	0.553
Partner Order (ABAB vs. ABBA)	0.03	0.03	0.83	0.408	0.05	0.03	1.65	0.100	0.12	0.16	0.73	0.465
Linguistic Co-presence * Visual Co-presence	-0.34	0.06	-5.41	<0.001	-0.40	0.06	-6.68	<0.001	-0.80	0.33	-2.44	0.015
Partner * Partner Order	0.07	0.06	1.06	0.288	0.16	0.06	2.72	0.007	-0.03	0.33	-0.10	0.919
Linguistic Co-presence * Partner	-0.08	0.06	-1.23	0.219	-0.03	0.06	-0.42	0.678	-0.34	0.33	-1.05	0.294
Visual Co-presence * Partner	0.20	0.06	3.10	<0.001	0.11	0.06	1.81	0.071	0.96	0.33	2.93	<0.01

Note. Each dependent measure is modeled as a function of the centered and contrast coded predictors Linguistic Co-Presence (Yes = -.5, No = .5), Visual Co-Presence (Yes = -.5, No = .5), Partner Order (ABBA = -.5 vs. ABAB = .5), Partner identity (Partner: B = -.5 vs. A = .5), and the interactions of these predictors, using the maximal random effect structure possible. For each fixed effect and interaction, we report the unstandardized coefficient and its standard error, along with the associated t-value and p-value, and the overall variance captured by the model (R²). Statistically significant predictors (at the Bonferroni adjusted $p = .0125$ level) are in bold. The variance and standard deviation of each random effect are reported in Table C1 in Appendix C.

Table 5

Mixed-effects models for the presence of hedges, the presence of definiteness (definite expressions and meta-references), and the directors' description onset times for their initial descriptions in Phase 2.

Predictor	Hedges				Definiteness				Description onset times			
	R ² = .39				R ² = .38				R ² = .40			
	B	SE	z	p	B	SE	z	p	B	SE	t	p
Intercept	1.69	0.24	-7.03	<0.001	-1.23	0.24	-5.07	<0.001	7.97	0.06	129.38	<0.001
Linguistic Copresence (Yes vs. No)	0.91	0.22	4.20	<0.001	-1.12	0.22	-4.99	<0.001	0.04	0.03	1.58	0.115
Visual Copresence (Yes vs. No)	0.08	0.26	-0.32	0.750	-0.08	0.19	-0.43	0.668	0.03	0.03	1.19	0.235
Partner (A vs. B)	0.16	0.18	-0.86	0.390	-0.00	0.19	-0.03	0.980	0.03	0.03	-1.18	0.238
Partner Order (ABAB vs. ABBA)	0.23	0.17	1.35	0.176	0.22	0.24	0.92	0.360	0.08	0.03	2.80	0.005
Linguistic Copresence * Visual Copresence	0.76	0.35	-2.17	0.030	-0.59	0.32	-1.83	0.068	0.17	0.05	-3.17	0.002
Partner * Partner Order	0.93	0.34	2.71	0.007	-0.28	0.32	-0.89	0.376	0.33	0.05	6.23	<0.001
Linguistic Copresence * Partner	0.25	0.35	0.72	0.471	-0.28	0.32	-0.87	0.382	0.06	0.05	1.03	0.303
Visual Copresence * Partner	0.09	0.34	0.25	0.800	0.03	0.32	0.10	0.920	0.01	0.05	0.23	0.822

Note. Each dependent measure is modeled as a function of the centered and contrast coded predictors Linguistic Co-Presence (Yes = -.5, No = .5), Visual Co-Presence (Yes = -.5, No = .5), Partner Order (ABBA = -.5 vs. ABAB = .5), Partner identity (Partner: B = -.5 vs. A = .5), and the interactions of these predictors, using the maximal random effect structure possible. For each fixed effect and interaction, we report the unstandardized coefficient and its standard error, along with the associated t-value or z-value, its p-value, and the overall variance captured by the model (R²). Statistically significant predictors (at the Bonferroni adjusted $p = .0125$ level) are in bold. The variance and standard deviation of each random effect are reported in Table C1 in Appendix C.

Provisionality: Hedges

Next, we considered the presence of hedges, taken to reflect the provisionality with which speakers describe a referent to their partner. Similar to words, idea units, and reconceptualizations, the conditions of co-presence in Phase 1 also impacted the directors' use of hedging in their initial descriptions in Phase 2. As shown in Table 5, directors were less likely to use hedges when referring to previously discussed items: linguistic co-presence was a significant predictor in the model. In addition, the effect of linguistic co-presence was accentuated by visual co-presence. Compared to referring to fully shared items, directors were 1.5 times more likely to use hedges when referring to items shared linguistically-only and 2.4 times more likely to use hedges when describing items shared visually-only (see Table 3). Nevertheless, the interaction of linguistic and visual co-presence was not a significant predictor of hedges ($p = .03$) under our conservative Bonferroni-corrected alpha level. As with the measures of content, visual co-presence was not a significant predictor of hedges on its own.

The incidence of hedges was also predicted by an interaction between partner identity and partner order. Directors were less likely to produce hedges with matcher A in order ABBA ($M = .15$, $SD = .36$) than in order ABAB ($M = .24$, $SD = .43$), and this difference across orders was greater than for Matcher B (for whom the incidence of hedges in order ABAB was: $M = .19$, $SD = .39$; in order ABBA: $M = .22$, $SD = .42$).

Definiteness: Definite Expressions and Metareferences

Definite expressions and metareferences, taken to reflect the marking of references as previously shared, showed sensitivity primarily to linguistic co-presence. In contrast to the measures capturing content, the interaction between linguistic and visual co-presence was not significant ($p = .07$; see Table 5). Instead, what mattered was whether items were previously discussed or

not: only the linguistic co-presence of items in Phase 1 significantly predicted the use of definiteness in Phase 2. Directors used definite expressions and meta-references in Phase 2 more frequently for items that had been linguistically shared with that particular partner in Phase 1 (whether shared fully or linguistically-only) than for items not linguistically shared (see Table 3).

Strikingly, for a given item, directors adapted their use of definiteness appropriately to each partner. Recall that items that were fully shared with one matcher were not shared at all with the other matcher, and items that were shared linguistically-only with one matcher were shared visually-only with the other. Figure 3 illustrates the use of definiteness in Phase 2 in a manner that reflects the complementary copresence conditions experienced within-pairs in Phase 1; we therefore juxtapose side-by-side LV vs. N items, and L-only vs. V-only items. This presentation underscores the feat of memory underlying directors' use of definiteness in Phase 2, with each director interacting with their two matchers. In order to use definiteness appropriately, directors were faced with recalling whether each tangram had been discussed or not with each matcher in Phase 1. They had to distinguish whether an item was shared fully with one partner and not at all with the other, and whether an item was shared only-linguistically with one partner and only-visually with the other. As Figure 3 illustrates, directors achieved this feat: they successfully distinguished previously discussed and non-discussed items between the two matchers, and they did so equally well for Matchers A and B.

Planning: Description onset time

Finally, we examined description onset latencies, taken to reflect speakers' planning of referring expressions (including any consideration of the co-presence status of a referent). In contrast to all previous measures, linguistic co-presence on its own did not predict description onset times. Instead, the interaction between linguistic and visual co-presence did. This is

because, as shown in Table 3, directors began describing tangrams that had been previously shared fully (both linguistically and visually) more quickly than in all other conditions.

There was also a significant interaction between partner identity and partner order. Directors began their tangram descriptions faster with Matcher A in order ABBA ($M = 2807.76$, $SD = 1342.13$) than ABAB ($M = 3654.00$, $SD = 2050.91$), consistent with the pattern observed with words and hedges. These interaction between partner identity and partner order for description onset times, words (see Table 4), and hedges (Table 5) occurred regardless of initial co-presence conditions, so are likely due to practice or fatigue.

[INSERT FIGURE 3 ABOUT HERE]

Figure 3

Figure 3 shows bar plots representing the mean proportion of items across the four co-presence conditions and across partner identity. Error bars represent the standard error of the mean. The ordering of the co-presence conditions reflects the design experienced within-pairs in Phase 1, illustrating the contrasts in definiteness that the speaker made when referring in Phase 2. This captures the complementary co-presence status of items: items that were LV with one partner were N with the other; items that were L-only with one partner were V-only with the other. To describe items appropriately, directors had to distinguish (as illustrated in the grey thought bubbles) that they had shared a given item fully with one partner and not at all with the other, or only-linguistically with one partner and only-visually with the other. As shown, directors distinguished items that had been discussed (LV or L-only) vs. non-discussed (V-only or N) with the each of the two matchers. These distinctions in definiteness between discussed and non-discussed items in our design (i.e., between LV vs. N, and between L-only vs. V-only) are illustrated in the arrows.

Summary

Speakers adapted the content of their descriptions in a gradient manner: they used fewer idea units and words when describing items that had been previously described only linguistically than only visually, and fewer still for items that had been shared fully. Linguistic co-presence had an independent effect on these measures, with descriptions being more attenuated and less provisional than when items were not previously discussed. In contrast, visual co-presence did *not* have an independent effect on descriptions. Nevertheless, as suggested by the gradient pattern observed and the interaction of the two constraints in the models, the effect of linguistic co-presence was amplified by visual co-presence for words and idea units. The same numerical patterns were also found for reconceptualizations and hedges, with perspective-switches and the provisionality of descriptions decreasing in the same gradient manner across co-presence conditions.

The speakers' use of definite expressions and their descriptions' onset times exhibited distinct patterns, consistent with audience design. Definite expressions were used in Phase 2 only for items previously discussed (i.e., with linguistic co-presence in Phase 1); neither visual co-presence nor its interaction with linguistic co-presence predicted definiteness. Finally, description onset times depended on the interaction of the two co-presence constraints, with speakers starting descriptions more quickly when items were fully shared compared to all other conditions.

Altogether, the patterns of sensitivity of the six measures to the two co-presence constraints suggest that speakers kept track of how they had previously shared information with their partners by adapting their descriptions appropriately when re-referring. Speakers distinguished unimodally and multimodally shared information in a gradient manner through

their choices of content, with linguistic experience weighted more heavily; they distinguished discussed vs. non-discussed items through their use of definiteness; and they distinguished fully shared items from all other items by starting their descriptions more readily. Directors made these distinctions despite the challenge of having shared each item in a complementary fashion across the two partners, demonstrating that they tracked co-presence in a partner-specific way.

Effects of initial co-presence (Phase 1) on speaker-addressee coordination (Phase 2)

Because referring depends on both speakers and addressees, we considered how successfully pairs coordinated in Phase 2 according to the number of turns over which they discussed cards and the number of errors they made. Overall, pairs were most efficient and most accurate in Phase 2 for items that had been fully shared in Phase 1. The linguistic co-presence, visual co-presence, and interaction of these two constraints during initial encoding in Phase 1 all predicted the number of conversational turns in Phase 2 (see Table 6), with items that had been fully shared items needing the fewest number of conversational turns in Phase 2 (Table 3). And in terms of task accuracy, pairs committed marginally fewer errors in Phase 2 ($p = .0198$; see Tables 3 and 6) when tangrams had been fully shared than in all other conditions (no other parameters in that model were significant).

Thus, both linguistic co-presence and visual co-presence influenced the efficiency of pairs in Phase 2, and they did so with a compounded effect. Fully sharing an item in Phase 1—in which the matcher ratified the identity of the referent that the director described by sharing it visually on the board—provided the strongest evidence for grounding. This episodic evidence later facilitated that pair's coordination during re-referring in Phase 2, with interaction taking the fewest turns and leading to numerically the fewest errors.

[INSERT TABLE 6 ABOUT HERE]

Table 6

Mixed-effects models for the number of turns (log-transformed) and errors in Phase 2.

Predictor	Turns				Errors			
	R ² = .32				R ² = .20			
	B	SE	t	p	B	SE	t	p
Intercept	1.19	0.06	21.12	<0.001	0.03	0.01	4.93	<0.001
Linguistic Copresence (Yes vs. No)	0.14	0.05	2.95	0.009	0.00	0.01	0.15	0.881
Visual Copresence (Yes vs. No)	0.12	0.03	3.97	<0.001	0.01	0.01	1.11	0.269
Partner (A vs. B)	-0.02	0.03	-0.83	0.408	0.01	0.01	0.79	0.432
Partner Order (ABAB vs. ABBA)	0.06	0.03	2.09	0.037	0.01	0.01	0.45	0.654
Linguistic Copresence * Visual Copresence	-0.29	0.06	-4.94	<0.001	-0.06	0.03	-2.34	0.020
Partner * Partner Order	0.15	0.06	2.45	0.015	0.02	0.03	0.77	0.442
Linguistic Copresence * Partner	-0.13	0.06	-2.12	0.035	-0.04	0.03	-1.41	0.159
Visual Copresence * Partner	0.07	0.06	1.16	0.248	0.02	0.03	0.79	0.432

Note. Each dependent measure is modeled as a function of the centered and contrast coded predictors Linguistic Co-Presence (Yes = -.5, No = .5), Visual Co-Presence (Yes = -.5, No = .5), Partner Order (ABBA = -.5 vs. ABAB = .5), Partner identity (Partner: B = -.5 vs. A = .5), and the interactions of these predictors, using the maximal random effect structure possible. For each fixed effect and interaction, we report the unstandardized coefficient and its standard error, along with the associated t-value, p-value, and the overall variance captured by the model (R²). Statistically significant predictors (at the Bonferroni adjusted $p = .0125$ level) are in bold. The variance and standard deviation of each random effect are reported in Table C2 in Appendix C.

Discussion

These findings show that in a spontaneous communicative setting, the distinct co-presence conditions under which a speaker first shared information with each of two different conversational partners shaped the referring expressions subsequently addressed to each partner. The speaker appeared to track how information had been shared with each partner, despite the fact that the initial shared experiences with each partner were subject to interference due to a high degree of similarity (e.g., involving the same objects).

Upon re-referring, speakers distinguished information that had been shared fully, only linguistically, or only visually, in a gradient manner, through the numbers of words and idea units in their initial descriptions. Reconceptualizations and hedges showed the same pattern numerically. Moreover, speakers also used definiteness in re-referring in a manner that distinguished a referent previously discussed with one partner from the same referent that had not been discussed before with the other. Further, their description onset times distinguished referents that had been fully shared (both linguistically and visually) from referents shared in only one modality or not at all: speakers began describing fully shared items more quickly than all other conditions. Linguistic and visual co-presence also influenced pairs' joint performance, which was most efficient when information had been shared fully.

This work goes beyond prior studies that have demonstrated that speakers adapt their descriptions in a partner-specific manner by distinguishing information that is new vs. old to a given partner (Brennan & Clark, 1996; Horton & Gerrig, 2005b; Galati & Brennan, 2010; 2014). The most novel contribution of the present study is that it demonstrates that speakers not only distinguish which information is new vs. old for each of the two partners, but also distinguish *how* the old information was shared with each partner. We take our findings to suggest that: (1)

Speakers can retain (at least briefly) memory traces that differ based on the modality through which they previously shared information with their addressees; (2) these traces are sufficient to shape speakers' subsequent referring expressions, even when such traces differ for different addressees of the same information; (3) subsequent referring expressions reflect appropriate description strategies consistent with the prior opportunities for grounding that information; and (4) despite these appropriate description strategies (signaling the status of shared information), joint task performance is affected by whether pairs could previously ground that information fully. We highlight and expand on the implications of these four findings for models of dialogue.

(1) Perceptual co-presence can leave traces in memory

Speakers' adaptations indicate that they do, in fact, retain in memory episodic information from the co-presence conditions for each item with each partner, long enough to successfully cue that information during subsequent utterance planning within a task episode. The clearest evidence comes from how speakers adapted the content of their initial referring expressions in Phase 2. This evidence suggests that speakers tracked both the linguistic and the visual co-presence status of each item *with that partner* during Phase 1. This is illustrated in the consistent patterns of idea units, words, and reconceptualizations, which show a gradient adaptation across unimodally and multimodally shared items. Speakers distinguished fully shared referents from referents shared only linguistically, as well as referents shared only linguistically from those shared only visually, such that numbers of idea units, words, reconceptualizations, and hedges in re-referring fell along the gradient: full co-presence < linguistic co-presence < visual co-presence. These findings are consistent with previous proposals that representations of common ground can be rich enough to permit gradient adaptation (as proposed by Brown-Schmidt, 2012), rather than tracking just an all-or-nothing

distinction of whether information was shared at all. Gradient adaptation makes sense if common ground is represented as several simple constraints (such as linguistic and visual co-presence) that can have compounded effects and can be weighted differently depending on context. In a referential communication task such as the one used here, referents were hard to lexicalize, so speakers may have prioritized linguistic co-presence over visual co-presence. Here, linguistic co-presence influenced descriptions during re-referring more than visual co-presence did, with speakers producing descriptions that were more definite, more concise, and less provisional when referring to previously discussed items. In another context, speakers might instead prioritize visual co-presence, as for example when collaborators must identify the spatial details of an event they have witnessed together. If the relative weighting of co-presence constraints is context-specific, depending on the goals and demands of the task, this would be consistent with source monitoring, where relevant weighted factors contribute to source assessment (Johnson, Hashtroudi, & Lindsay, 1993).

(2) Co-presence conditions can have compounded effects during re-referring

Although speakers prioritized linguistic over visual co-presence when formulating their descriptions, they didn't encode and track linguistic co-presence alone. We found clear evidence that speakers retained both linguistic and visual co-presence in memory and that the effects of these constraints were compounded for some aspects of re-referring. The effect of linguistic co-presence was amplified by visual co-presence, as indicated by the interaction of the two constraints (with less content in terms of words and idea units, as well as faster onset times for referents that had been shared both linguistically and visually).

What kind of model best captures these findings? The interaction of linguistic and visual constraints is well accommodated by constraint-based models (Hanna, Tanenhaus, & Trueswell,

2003; Brown-Schmidt & Hanna, 2011; Tanenhaus & Trueswell, 1995) and dynamical systems models (Duran & Dale, 2014; Dale et al., 2018) that can represent multiple constraints or control parameter. Neural network models can bridge these two approaches (Duran et al., 2016). A resonance-based model is compatible as well, where multiple, previously-encoded cues act in parallel until their resonance reaches threshold and can influence other processes (Horton & Brennan, 2016; Horton & Gerrig, 2016; Ratcliff, 1978; Ratcliff & McKoon, 1988). These approaches are consistent with ordinary memory processes, yet can capture “rich” representations that support adaptive behavior in dialogue and perspective-taking.

The interactive alignment model (Pickering & Garrod, 2004), as well as models that argue that language use is initially “egocentric” (with any partner-specific adjustments made later; e.g., Barr & Keysar, 2006; Horton & Keysar, 1996; Keysar, Barr, & Horton, 1998; Keysar, Barr, Balin, & Paek, 1998; Kronmüller & Barr, 2007) is incompatible with our partner-specific findings. Such models have focused mainly on priming of linguistic forms from the speaker’s own experience or perspective, with an emphasis on the activation of recent experiences and without a role for visual experiences. We found that new linguistic forms in Phase 2 (i.e., reconceptualizations) were less frequent when an item had been discussed with that partner in Phase 1. Although this finding can be seen as consistent with the interactive alignment model and egocentric explanations (insofar as fewer reconceptualizations imply more reused idea units), the fact that the effect was partner-specific, and that it was amplified when the item had also been shared visually, cannot.

Although our experiment was not designed to test attenuation, there may also have been some shortening across the descriptions of Phase 1 and Phase 2 that was driven by the speaker’s own experience. This possibility would be consistent with accounts that model both egocentric

and other-centric factors as constraints that shape language production and interpretation (Brown-Schmidt & Hanna, 2011; Galati & Brennan, 2014; Heller et al., 2016; Mozuraitis et al., 2018). There was some limited suggestive evidence of that: some measures (words and hedges) indicated an interaction between partner identity and partner order (with more overall attenuation for partner A in order ABBA). However, because this interaction effect does not reflect the role of modality in Phase 1, it suggests a general effect of practice or fatigue. Other findings further underscore that priming alone isn't an explanation. Speakers didn't simply rely on their most recent linguistic experience, which should have yielded stronger effects of linguistic co-presence for the more recent partner from Phase 1. We did not find any evidence of an interaction between linguistic co-presence and partner identity: speakers did not attenuate their descriptions any more with the more recent partner (B) than with the other partner (A).

Altogether, neither the interactive alignment model nor for-the-speaker (egocentric) priming explain our finding that the descriptions in Phase 2 were partner-specific *and* sensitive to how information had been shared (including visually). Over and above any practice or fatigue effects based on their own experiences, speakers tailored their descriptions to each partner according to how items were previously shared: adaptation was both partner-specific and modality-specific.

(3) Co-presence conditions influence the selection of description strategies

When considered as a whole, speakers' linguistic adjustments reflect appropriate strategies for signaling the information status of items. Our findings from Phase 1—that linguistic-only copresence requires more effort (according to multiple measures) than multimodal or “full” co-presence—suggest that information that had been shared multimodally was grounded to a greater degree (with stronger perceptual evidence) than unimodally shared information. The

process of grounding is shaped by communicative affordances, including whether speakers have visual evidence of their partners' understanding (Brennan, 2005; Clark & Brennan, 1991; Clark & Krych, 2004; Gergle, Kraut & Fussell, 2004). Keeping track of the strength of evidence in grounding—which is strongly affected by the conditions of co-presence under which information has been shared—is relevant in a host of collaborative tasks, as it permits interlocutors to gauge how much effort they need to expend to reach an appropriate grounding criterion for current purposes (Clark & Wilkes-Gibbs, 1986; Clark & Schaefer, 1989).

As a consequence of prior grounding, re-referring in Phase 2 raised different kinds of considerations for multimodally shared and unimodally shared items. Re-referring to tangrams that had been partially shared (only linguistically or only visually) required speakers to determine how to best package referring expressions, with an increase in planning time relative to the more straightforward case of re-referring to fully shared items; speakers began describing fully shared tangrams more quickly than tangrams in all the other conditions. Although this could mean that speakers failed to distinguish initially between referents that were shared partially or not at all, we take this to mean that each of these conditions was associated with a processing cost for selecting an appropriate description strategy.

Speakers' ability to use definite expressions appropriately in Phase 2 with each partner was particularly impressive. Definite expressions were used to mark prior linguistic co-presence status (i.e., that an item had been discussed); they were used similarly often for items for which there had been full or linguistic-only co-presence. At the same time, speakers differentiated between items shared fully or linguistically-only by attenuating or expanding upon prior descriptions, signaling the strength of evidence of their prior grounding. This finding on the part of speakers complements evidence that addressees, as well, keep track of whether they have

completed grounding a referent (as opposed to when that process was interrupted; Brown-Schmidt, 2009). In Brown-Schmidt's study, addressees distinguished items that had been “fully grounded” from those that had not. Upon hearing a temporarily ambiguous question (e.g., “What’s above the cow that’s wearing shoes?” in a context with two cows, one with shoes and one with glasses), listeners were more likely to direct their attention, as indicated by their eye movements, away from a competitor that had been previously mentioned and fully grounded (i.e., an animal above one of the two cows previously mentioned) and toward the unmentioned target (an animal above the other cow that was not previously mentioned). However, when the competitor had been mentioned but had not been fully grounded due to an interruption in the exchange before the partner could respond, the effect was attenuated. This is converging evidence that conversational partners are capable of tracking not only whether they have previously discussed something, but also the strength of evidence used to ground it.

In our study, the lengths of the interactions varied according to the grounding opportunities afforded in the different co-presence conditions. This raises the possibility that, instead of (or in addition to) tracking details about modality, speakers recalled having interacted longer over tangrams shared only linguistically (which took longer on average than fully-shared items), and/or recalled their addressees being unable to find tangrams shared only linguistically (for addressees, this condition was like a target-absent visual search). Because the length of the interaction could well affect an episodic memory, it could potentially serve as a cue to reconstruct the co-presence status of the shared information. Consistent with the “persons-as-contexts” account of partner-specific adaptation (Brown-Schmidt, Yoon, & Ryskin 2015), speakers can use their conversational partner’s identity as a contextual cue to retrieve associations from memory relevant to their shared experience (see also Horton & Brennan, 2016;

Horton & Gerrig, 2002; 2005a; 2005b; 2016; Metzing & Brennan, 2003). That is, the partner's identity cues an episode that tags the co-presence status of a referent (e.g., in our study, recalling that the partner placed a joker card on the board, perhaps saying "I don't have it"), or a more extended episodic experience of that referent with that partner (i.e., the length of the interaction), or else other affective aspects of that experience (e.g., that it was annoying to match that item). We acknowledge that in our study, such temporal and affective factors covary with the referent's co-presence status and may be part of the resonating traces in memory for that experience. With this caveat, we frame speakers' partner-specific adaptation in terms of co-presence, as that was systematically manipulated in Phase 1. Future studies that disentangle co-presence constraints from covarying factors (e.g., by imposing time limits to manipulate the length of the interaction) may shed light on the specific episodic information that speakers draw from as they tailor their descriptions to their conversational partners.

Just as items shared only linguistically led to distinctive description profiles in Phase 2, so did items shared only visually. The latter cases led to the shortest episodes in Phase 1, as directors simply placed their two leftover cards on the board. In Phase 2, speakers marked those items as newer (by using more idea units, words, and reconceptualizations) than items shared only linguistically, and in ways similar to completely new referents (e.g., using similar amounts of definite expressions). However, they treated items shared only visually as even "newer" than referents that were actually new to a given partner: namely, by using numerically more idea units, words, and reconceptualizations. This may be because, to determine whether an item has been previously shared visually with a particular partner, speakers must access a memory trace for that experience, whereas when faced with a new item for a particular partner, speakers must infer its status in the absence of perceptual evidence, upon failing to access any episodic trace for

that partner-item association (recall that items new to one partner had been fully shared with the other partner). Even though the experience of sharing items visually-only was brief, being able to access a memory trace for that experience with a partner may have served as stronger evidence for the information status of that item (and thereby, a need to provide an expanded description), compared to having to infer that an item must be new to a partner on the basis of the absence of a memory trace.

Given our design, the descriptions of items to one partner may have been affected by the co-presence status of those items with the other partner. That is, in Phase 2 when describing a V-only item, speakers may have recalled that it had been shared as L-only with the other partner and that it took many words to describe in Phase 1, suggesting it was a particularly difficult item—perhaps leading to a kind of overcompensation (resulting in even lengthier initial Phase 2 descriptions to the other partner). This is speculative, as there is no direct evidence for this. It is still possible that speakers expanded their descriptions of V-only items (compared to New, numerically speaking) because they could more readily establish how they shared those items with their current partner in Phase 1.

Despite substantial potential for interference in this task, each addressee's identity was remarkably successful in cueing the appropriate co-presence status. In most previous studies of audience design, the conversational partner's informational needs could be summarized as single, global constraint that could be applied over the entire interaction with that partner—for instance, for cuing category associations for subsets of items (e.g., *with this partner, I have matched cards of dogs, but not of turtles*, e.g., Horton & Gerrig, 2002, 2005b) or for cuing the status of information of a more extended stretch of discourse (e.g., *this partner has heard this entire story before*, e.g., Galati & Brennan, 2010, or *is knowledgeable about these objects*, e.g.,

Yoon & Brown-Schmidt, 2019). In the conversational context of Phase 2, speakers described individual items with the appropriate levels of informative detail or attenuation to specific partners, even though the items' information status varied along more than a single categorical constraint. In contrast, interactive alignment (Pickering & Garrod, 2004) would predict that all items should be similarly attenuated in Phase 2 insofar as they were all described by the speaker to one addressee or the other in Phase 1.

Although our data don't speak to the specific content of speakers' rich episodic memories of shared experiences with their partners, our findings suggest that modeling linguistic and visual co-presence as multiple binary constraints in memory can account for partner-specific adaptation in our study. Cueing information about co-presence in order to deploy distinct description strategies reflects how this information was previously grounded.

(4) Effective coordination during re-referring depends on the previous grounding process

Speakers' co-presence with particular addressees shaped appropriate referential strategies in both phases of the task and in all co-presence conditions; however, this did not always lead to effective communication. The fact that pairs coordinated most effectively (with fewer conversational turns and numerically fewer errors) during Phase 2 on items shared both linguistically and visually in Phase 1 suggests that having had strong evidence for grounding allows addressees to benefit maximally from speakers' adaptations. Although a speaker's adaptation may be informative and may be shaped by their intentions toward a partner, it may not be processed effectively by the partner to recover that information (see Brennan & Williams, 1995, for discussion of the criteria needed for cues to succeed in being "communicative"). Some linguistic adaptations failed to be clear to the addressee, despite the speaker's presumed intentions, insofar they were not processed effectively by the partner and thus did not facilitate

coordination. For instance, sometimes matchers could not readily identify items that were previously shared only-linguistically even when speakers signaled they had been discussed before by using definite expressions (evidence for this claim can be found in our online supplement that reports the partners' use of feedback, including requests for clarification and expansion). Even if they don't always benefit from the linguistic cues that speakers provide for them, addressees can serve as cues for speakers about how information has been previously shared (see also Horton & Gerrig, 2005a; Brown-Schmidt, Yoon, & Ryskin, 2015).

Conclusion

We have shown that speakers are capable of retaining multiple distinctions in memory about the perceptual modalities through which they have shared information with their conversational partners; that these perceptual distinctions are weighted differently; and that the information can be used later to cue referential strategies appropriate for particular partners. Upon re-referring, speakers' descriptions in their first uninterrupted turn demonstrate that they can quickly cue the conditions of previous co-presence, especially for referents that had been fully shared using both visual and linguistic co-presence. Our earlier work (Galati & Brennan, 2010; 2014; Brennan, Galati, & Kuhlen, 2010) established that speakers can track simple, binary distinctions about a partner's knowledge. But the gradient adaptation documented here demonstrates that speakers can maintain richer representations about prior shared experiences that can shape utterance planning.

The tracking of multiple distinctions about common ground is well-accommodated by a constraint-based model of language processing, which allows for multiple constraints to be weighted depending on their availability and relevance, and to be integrated with each other in parallel and probabilistically (e.g., Hanna, Tanenhaus, & Trueswell, 2003). It can also be

accommodated by other formal models of adaptive behavior in dialogue, including dynamical ones that can represent a host of simple constraints and their interactions (Duran & Dale, 2014; Duran et al., 2016; Dale et al., 2018). The linguistic adaptation documented here cannot be explained by appealing only to an automatic mechanism like priming (e.g., Pickering & Garrod, 2004), as the effects of linguistic co-presence did not depend on recency, but on partner identity. Although there were some effects of practice for some measures, speakers adapted based on how they had previously shared a given item with a particular addressee.

In sum, speakers are able to retain *what* information they have shared with *whom*, and *under what conditions* they have shared it, at least long enough to support subsequent interactions with two different addressees. The co-presence conditions between speakers and addressees affect not only how speakers coordinate in the *current* conversation, but also what description strategies they are likely to use *later*. Distinctions in co-presence modality can be sufficiently simple to leave traces in memory and be retrieved rapidly, supporting the gradient use of common ground in utterance planning.

References

- Anderson, J. R. & Bower, G. H. (1973). *Human associative memory*. Washington, DC: Winston.
- Balota, D.A., Yap, M.J., Cortese, M.J., Hutchison, K.A., Kessler, B., Loftis, B., Neely, J.H., Nelson, D.L., Simpson, G.B., & Treiman, R. (2007). The English Lexicon Project. *Behavior Research Methods*, 39, 445-459.
- Bard, E. G., Anderson, A. H., Sotillo, C., Aylett, M., Doherty-Sneddon, G., & Newlands, A. (2000). Controlling the intelligibility of referring expressions in dialogue. *Journal of Memory & Language*, 42, 1-22.
- Bard, E. G., & Aylett, M. P. (2001). Referential form, word duration, and modeling the listener in spoken dialogue. Paper presented at the Twenty-third Annual Conference of the Cognitive Science Society, Edinburgh, Scotland.
- Barr, D. J., & Keysar, B. (2006). Perspective taking and the coordination of meaning in language use. In: Traxler, M.J. & Gernsbacher, M.A. (Eds.), *Handbook of Psycholinguistics*, (pp. 901-938). Academic Press/Elsevier. <https://doi.org/10.1016/B978-012369374-7/50024-9>
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255-278.
- Bates, D., Kliegl, R., Vasishth, S., & Baayen, H. (2015). Parsimonious mixed models. arXiv preprint arXiv:1506.04967.
- Bates, D., Mächler, M., Bolker, B. & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67, 1-48.
- Brennan, S. E. (2005). How conversation is shaped by visual and spoken evidence. In J. Trueswell & M. Tanenhaus (Eds.), *Approaches to studying world-situated language use:*

Bridging the language-as-product and language-action traditions (pp. 95-129).

Cambridge, MA: MIT Press.

Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 1482-1493.

Brennan, S. E., Galati, A., & Kuhlen, A. K. (2010). Two minds, one dialog: Coordinating speaking and understanding. In B. Ross (Ed.) *Psychology of Learning and Motivation*, 53, (pp. 301–345). Academic Press/Elsevier.

Brennan, S. E., & Hanna, J. E. (2009). Partner-specific adaptation in dialogue. *Topics in Cognitive Science (Special Issue on Joint Action)*, 1, 274-291.

Brennan, S. E., & Williams, M. (1995). The feeling of another's knowing: Prosody and filled pauses as cues to listeners about the metacognitive states of speakers. *Journal of Memory and Language*, 34(3), 383-398. <https://doi.org/10.1006/jmla.1995.1017>

Brennan, S. E., & Williams, M. (1995). The feeling of another's knowing: Prosody and filled pauses as cues to listeners about the metacognitive states of speakers. *Journal of Memory and Language*, 34, 383-398.

Brown, P., & Dell, G. S. (1987). Adapting production to comprehension: The explicit mention of instruments. *Cognitive Psychology*, 19, 441-472.

Brown-Schmidt, S. (2009). The role of executive function in perspective taking during online language comprehension. *Psychonomic Bulletin and Review*, 16 (5), 893-900.

Brown-Schmidt, S., & Hanna, J. E. (2011). Talking in another person's shoes: Incremental perspective-taking in language processing. *Dialog and Discourse*, 2, 11-33.

- Brown-Schmidt, S. (2012). Beyond common and privileged: Gradient representations of common ground in real-time language use, *Language and Cognitive Processes*, 27 (1), 62-89.
- Brown-Schmidt, S. & Horton, W. S. (2014). The Influence of Partner-Specific Memory Associations on Picture Naming: A Failure to Replicate Horton (2007). *PLoS ONE* 9(10): e109035. doi:10.1371/journal.pon e.01 09 035
- Brown-Schmidt, S., Yoon, S. O., & Ryskin, R. A. (2015). People as Contexts in Conversation. In B. Ross (Ed.), *The Psychology of Learning and Motivation* (pp. 59–99). Academic Press: Elsevier Inc.
- Brugman, H., & Russel, A. (2004). Annotating Multi-media / Multi-modal resources with ELAN. *LREC*.
- Butterfill, S. A., & Apperly, I. A. (2013). How to construct a minimal theory of mind. *Mind & Language*, 28(5), 606-637. <https://doi.org/10.1111/mila.12036>
- Carney, R. N., & Levin, J. R. (2002). Pictorial illustrations still improve students' learning from text. *Educational Psychology Review*, 13, 5-26.
- Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. In L. Resnick, J. Levine, & S. Teasley, (Eds). *Perspectives on socially shared cognition* (pp. 127-149). Washington, DC: APA Books.
- Clark, H. H., & Krych, M. (2004). Speaking while monitoring addressees for understanding, *Journal of Memory and Language*, 50, 62-81.
- Clark, H. H., & Marshall, C. E. (1981). Definite reference and mutual knowledge. In A. J. Jodhi, B. L. Webber, & I. A. Sag (Eds.), *Elements of discourse understanding*, (pp. 10-63). Cambridge: Cambridge University Press.

- Clark, H. H., & Schaefer, E. F. (1989). Contributing to discourse. *Cognitive Science*, 13, 259-294.
- Clark, H. H., & Wilkes-Gibbs, D. (1986). Referring as a collaborative process. *Cognition*, 22, 1-39.
- Cooper, W. E., & Paccia-Cooper, J. (1980). Syntax and speech. Cambridge, MA: Harvard University Press.
- Dale, R., Galati, A., Alviar, C., Contreras Kallens, P., Ramirez-Aristizabal, A. G., Tabatabaieian, M., & Vinson, D. W. (2018). Interacting Timescales in Perspective-Taking. *Frontiers in Psychology*, 9, 1278. <https://doi.org/10.3389/fpsyg.2018.01278>
- Duran, N. D., & Dale, R. (2014). Perspective-taking in dialogue as self-organization under social constraints. *New Ideas in Psychology*, 32, 131–146.
<https://doi.org/10.1016/j.newideapsych.2013.03.004>
- Duran, N., Dale, R., & Galati, A. (2016). Toward Integrative Dynamic Models for Adaptive Perspective Taking. *Topics in Cognitive Science*, 8(4), 761–779.
<https://doi.org/10.1111/tops.12219>
- Epley, N., Morewedge, C. K., & Keysar, B. (2004). Perspective taking in children and adults: Equivalent egocentrism but differential correction. *Journal of Experimental Social Psychology*, 40(6), 760-768. <https://doi.org/10.1016/j.jesp.2004.02.002>
- Galati, A., & Avraamides, M. N. (2015). Social and representational cues jointly influence spatial perspective-taking. *Cognitive Science*, 39(4), 739-765.
<https://doi.org/10.1111/cogs.12173>

- Galati, A., Michael, C., Mello, C., Greenauer, N. M., & Avraamides, M. N. (2013). The conversational partner's perspective affects spatial memory and descriptions. *Journal of Memory and Language*, 68(2), 140-159. <https://doi.org/10.1016/j.jml.2012.10.001>
- Galati, A., & Brennan, S. E. (2010). Attenuating information in Spoken Communication: For the speaker, or for the addressee? *Journal of Memory and Language*, 62, 35-51.
- Galati, A., & Brennan, S. E. (2014). Speakers adapt gestures to addressees' knowledge: Implications for models of co-speech gesture. *Language, Cognition and Neuroscience*, 29(4), 435–451. <https://doi.org/10.1080/01690965.2013.796397>
- Gelman, A., Hill, J., & Yajima, M. (2012). Why we (usually) don't have to worry about multiple comparisons. *Journal of Research on Educational Effectiveness*, 5(2), 189-211.
- Gerrig, R. J., Brennan, S. E., & Ohaeri, J. O. (2000). What can we conclude from speakers behaving badly? *Discourse Processes*, 29, 173-178.
- Gergle, D., Kraut, R. E., Fussell, S. R. (2004). Language efficiency and visual technology: Minimizing collaborative effort with visual information. *Journal of Language and Social Psychology*, 23 (4), 491-517.
- Hanna, J. E., Tanenhaus, M. K., & Trueswell, J. C. (2003). The effects of common ground and perspective on domains of referential interpretation. *Journal of Memory and Language*, 49, 43–61.
- Hanna, J. E., & Tanenhaus, M. K. (2004). Pragmatic effects on reference resolution in a collaborative task: Evidence from eye movements. *Cognitive Science*, 28, 105–115.
- Heller, D., Parisien, C., & Stevenson, S. (2016). Perspective-taking behavior as the probabilistic weighing of multiple domains. *Cognition*, 149, 104–120. <https://doi.org/10.1016/j.cognition.2015.12.008>

- Hilliard, C., & Cook, S. W. (2016). Bridging gaps in common ground: Speakers design their gestures for their listeners. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(1), 91–103. <https://doi.org/10.1037/xlm0000154>
- Horton, W. S., (2007). The influence of partner-specific memory associations on language production: Evidence from picture naming. *Language and Cognitive Processes*, 22, 1114-1139.
- Horton, W. S., & Gerrig, R. J. (2005a). Conversational common ground and memory processes in language production. *Discourse Processes*, 40, 1-35.
- Horton, W. S., & Gerrig, R. J. (2005b). The impact of memory demands on audience design during language production. *Cognition*, 96, 127-142.
- Horton, W. S., & Gerrig, R. J. (2002). Speaker's experiences and audience design: knowing when and knowing how to adjust utterances to addressees. *Journal of Memory and Language*, 47, 589-606.
- Horton, W. S., & Gerrig, R. J. (2016). Revisiting the memory-based processing approach to common ground. *Topics in Cognitive Science*, 8, 780-795.
- Horton, W. S., & Keysar, B. (1996). When do speakers take into account common ground? *Cognition*, 59, 91-117.
- Hwang, J., Brennan, S. E., & Huffman, M. K. (2015). Phonetic adaptation in non-native spoken dialogue: Effects of priming and audience design. *Journal of Memory and Language*, 81, 72-90.
- Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin*, 114(1), 3-28.
- Johnson, P. (2014). Extension of Nakagawa & Schielzeth's R^2_{GLMM} to all random slopes models.

- Methods in Ecology and Evolution*, 5(9), 944–946.
- Keysar, B. (1997). Unconfounding common ground. *Discourse Processes*, 24, 253-270.
- Keysar, B., Barr, D. J., Balin, J. A., & Paek, T. S. (1998). Definite reference and mutual knowledge: process models of common ground in comprehension. *Journal of Memory and Language*, 39, 1–20.
- Keysar, B., Barr, D. J., & Horton, W. S. (1998). The egocentric bias of language use: Insights from a processing approach. *Current Directions in Psychological Science*, 7, 46-50.
- Kraljic, T. & Brennan, S.E. (2005). Prosodic disambiguation of syntactic structure: For the speaker or for the addressee? *Cognitive Psychology*, 50, 194-231.
- Kronmüller, E. & Barr, D. J. (2007). Perspective-free pragmatics: Broken precedents and the recovery-from-preemption hypothesis. *Journal of Memory and Language* 56 (2007) 436–455.
- Hayes, A. F., & Krippendorff, K. (2007). Answering the call for a standard reliability measure for coding data. *Communication Methods and Measures*, 1, 77-89.
- Horton, W. S., & Brennan, S. E. (2016). The role of metarepresentation in the production and resolution of referring expressions. *Frontiers in Psychology*, 7, 1111.
<https://doi.org/10.3389/fpsyg.2016.01111>
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2016). lmerTest: Tests in linear mixed effects models (R package version 2.0-33)[Computer software].
- Lakens, D. (2016, February 14). Why you don't need to adjust your alpha level for all tests you'll do in your lifetime. [Blog post]. Retrieved from
<http://daniellakens.blogspot.com/2016/02/why-you-dont-need-to-adjust-you-alpha.html>

- Lockridge, C. B., & Brennan, S. E. (2002). Addressees' needs influence speakers' early syntactic choices. *Psychonomic Bulletin and Review*, 9, 550-557.
- MacDonald, M. C. (1994). Probabilistic constraints and syntactic ambiguity resolution. *Language and Cognitive Processes*, 9, 157-201.
- Matthews, D., Lieven, E., & Tomasello, M. (2010). What's in a manner of speaking? Children's sensitivity to partner-specific referential precedents. *Developmental Psychology*, 46(4), 749-760.
- McRae, K., Spivey-Knowlton, M. J., & Tanenhaus, M. K. (1998). Modeling the influence of thematic fit (and other constraints) in on-line sentence comprehension. *Journal of Memory and Language*, 38, 283-312.
- Metzing, C. & Brennan, S. E. (2003). When conceptual pacts are broken: Partner-specific effects in the comprehension of referring expressions. *Journal of Memory and Language*, 49, 201-213.
- Mozuraitis, M., Stevenson, S., & Heller, D. (2018). Modeling Reference Production as the Probabilistic Combination of Multiple Perspectives. *Cognitive Science*, 42, 974-1008.
<https://doi.org/10.1111/cogs.12582>
- Murray, M. M., Michel, C.M., Grave de Peralta Menendez, R., Ortigue, S., Brunet, D., Gonzalez Andino, S. L., & Schnider, A. (2004). Rapid discrimination of visual and multisensory memories revealed by electrical neuroimaging, *Neuroimage*, 21 (1), 125-135.
<https://doi.org/10.1016/j.neuroimage.2003.09.035>
- Nadig, A. S., & Sedivy, J. S. (2002). Evidence of perspective-taking constraints in children's on-line reference resolution, *Psychological Science*, 13, 329-336.

- Paivio, A. (1986). *Mental representations: A dual coding approach*. Oxford, England: Oxford University Press.
- Pickering, M. J., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27, 167-226.
- Pickering, M. J., & Garrod, S. (2013). An integrated theory of language production and comprehension. *Behavioral and Brain Sciences*, 27, 167-226.
- R Core Team (2016). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59-108.
<https://doi.org/10.1037/0033-295X.85.2.59>
- Ratcliff, R., & McKoon, G. (1988). A retrieval theory of priming in memory. *Psychological Review*, 95(3), 385-408.
- Ryskin, R. A., Brown-Schmidt, S., Canseco-Gonzalez, E., Yiu, L. K., & Nguyen, E. T. (2014). Visuospatial perspective-taking in conversation and the role of bilingual experience. *Journal of Memory and Language*, 74, 46-76.
- Ryskin, R. A., Wang, R. F., & Brown-Schmidt, S. (2016). Listeners use speaker identity to access representations of spatial perspective during online language comprehension. *Cognition*, 147, 75-84.
- Sachs, J. S. (1967). Recognition memory for syntactic and semantic aspects of connected discourse. *Perception & Psychophysics*, 2(9), 437-442.
- Schober, M.F. (1993). Spatial perspective-taking in conversation. *Cognition*, 47, 1-24.
- Shams, L., & Seitz, A., (2008). Benefits of multisensory learning. *Trends in Cognitive Science*, 12, 411-417.

- Shintel, H., & Keysar, B. (2009). Less is more: A minimalist account of joint action in communication. *Topics in Cognitive Science*, 1(2), 260-273.
<https://doi.org/10.1111/j.1756-8765.2009.01018.x>
- Snodgrass, J. G., & Vanderward, M. (1980). A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning and Memory*, 6, 174-215.
- Stalnaker, R. (1974). Pragmatic Presuppositions. In M. K. Munitz & P. K. Unger (eds.), *Semantics and Philosophy* (pp. 197-213). New York University Press, New York.
University Press
- Swets, B., Jacovina, M. E., & Gerrig, R. J. (2013). Effects of conversational pressures on speech planning. *Discourse Processes*, 50, 23-51
- Szekely, A., D'Amico, S., Devescovi, A., Federmeier, K., Herron, D., Iyer, G., Jacobsen, T., & Bates, E. (2003) Timed picture naming: Extended norms and validation against previous studies, *Behavior Research Methods, Instruments & Computers*, 35 (4), 621-633.
- Tanenhaus, M. K., & Trueswell, J. C. (1995). Sentence comprehension. In J. Miller, & P. Eimas (Eds.), *Handbook of cognition and perception*. San Diego, CA: Academic Press.
- Yoon, S. O., Koh, S., & Brown-Schmidt. (2012). Influence of perspective and goals on reference production in conversation. *Psychonomic Bulletin & Review*, 19, 699–707.
- Yoon, S. O., & Brown-Schmidt. (2018). Aim Low: Mechanisms of Audience Design in Multiparty Conversation, *Discourse Processes*, 55 (7), 566–592.
- Yoon, S. O., & Brown-Schmidt. (2019). Contextual Integration in Multiparty Audience Design. *Cognitive Science*, 43(12). DOI: 10.1111/cogs.12807

Appendix A

Example of coding idea units (in Phase 1 and 2) and reconceptualizations (in Phase 2). Director 30 describes item 1.8 to Matcher A in Phase 1, and to both matchers in Phase 2.

According to the transcription conventions, pauses are marked /. Fillers are annotated as <uh> or <um>. The end point of a self-interruption is marked with an asterisk (*), and curly brackets {} contain undecipherable speech. Metacomments are in italics, and are not taken to contribute to content. Instrumental actions are in square brackets [].

To code for reconceptualizations, the idea units of the initial description in a given Phase 2 round are first determined, and then compared to the total idea units of a description of the same item in Phase 1. The comparison involves identifying which idea units in Phase 2 overlap with Phase 1, which are unique to Phase 2, and which are unique to Phase 1. Reconceptualizations are the idea units that are unique to Phase 2.

Round 1 (Phase 1) to Matcher A:

D30: okay
 this looks like /
 a<a> / ballerina /
 on a triangle*
 upside down triangle / stage /
 <um> it's doing a pirouette /
or I don't know what those things are /
 %laugh
 M_A: %laugh
 I don't think I have it /
 [puts up joker card]
 D30: *even if I knew {I...}*
 %laugh
 M_A: %laugh

Idea units of initial description (7): 1. ballerina 2. on triangle 3. upside down (triangle stage) 4. (upside down) triangle (stage) 5. (upside down triangle) stage 6. doing (pirouette) 7. (doing) pirouette

Idea units of total description (7): 1. ballerina 2. on triangle 3. upside down (triangle stage) 4. (upside down) triangle (stage) 5. (upside down triangle) stage 6. doing (pirouette) 7. (doing) pirouette

Round 3 (Phase 2) to Matcher B:

D30: okay
 this look like
 somebody / <um>
 twirling on a st* <uh>
 triangle s*

upside | down triangle stage ? /
 with one leg up in the air | /
 {on} pointy toes /
 M_B: okay /
 this<s> ? /
 [puts up standing man card]
 D30: yep

Idea units of initial description (11): 1. somebody 2. twirling 3. upside down (triangle stage) 4. (upside down) triangle (stage) 5. (upside down triangle) stage 6. one (leg) 7. (one) leg 8. (one leg) up (in the air) 9. (one leg up) in the air 10. pointy (toes) 11. (pointy) toes

Idea units overlapping between Phase 1 and Phase 2 (3): 1. upside down (triangle stage) = upside down (triangle stage) 2. (upside down) triangle (stage) = (upside down) triangle (stage) 3. (upside down triangle) stage = (upside down triangle) stage

Idea units unique to the initial description of Phase 2 (Reconceptualizations) (8): 1. somebody 2. twirling 3. one (leg) 4. (one) leg 5. (one leg) up (in the air) 6. (one leg up) in the air 7. pointy (toes) 8. (pointy) toes

Idea units unique to Phase 1 (4): 1. ballerina 2. on triangle 3. doing (pirouette) 4. (doing) pirouette

Round 4 (Phase 2) to Matcher A:

D30: this is the ballerina /
 %laugh
 M_A: %laugh
 [puts up standing man card]
 D30: on the stage /

Idea units of initial description (1): 1. ballerina

Idea units overlapping between Phase 1 and Phase 2 (1): 1. ballerina

Idea units unique to the initial description of Phase 2 (Reconceptualizations) (0): 0

Idea units unique to Phase 1 (6): 1. on triangle 2. upside down (triangle stage) 3. (upside down) triangle (stage) 4. (upside down triangle) stage 5. doing (pirouette) 6. (doing) pirouette

Appendix B

Table B1

Mixed-effects models for total words (log-transformed), total idea units (log-transformed), the proportion of items with hedges, number of turns, and number of errors in Phase 1.

Predictor	Total idea units				Total words				Items with hedges			
	R ² = .52				R ² = .52				R ² = .39			
	B	SE	t	p	B	SE	t	p	B	SE	z	p
Intercept	2.25	0.07	30.93	<0.001	3.63	0.07	49.06	<0.001	0.16	0.29	0.55	0.580
Co-Presence: L vs. LV	0.55	0.07	7.71	<0.001	0.60	0.08	7.02	<0.001	0.69	0.26	2.71	0.007
Partner: A vs. B	0.03	0.06	0.57	0.575	0.06	0.07	0.87	0.394	0.05	0.21	0.24	0.808
Co-Presence: L vs. LV * Partner: A vs. B	0.34	0.08	4.03	<0.001	0.28	0.09	3.11	0.002	-0.35	0.42	-0.85	0.394
Predictor	Number of turns				Number of errors							
	R ² = .36				R ² = .19							
	B	SE	t	p	B	SE	t	p				
Intercept	6.09	0.46	13.21	<0.001	0.12	0.03	4.43	<0.001				
Co-Presence: L vs. LV	2.45	0.69	3.56	0.002	0.15	0.05	3.21	0.004				
Partner: A vs. B	0.60	0.63	0.96	0.348	0.05	0.04	1.35	0.190				
Co-Presence: L vs. LV * Partner: A vs. B	1.46	0.72	2.02	0.044	0.02	0.06	0.25	0.805				

Note. Each dependent measure is modeled as a function of the centered and contrast coded predictors: condition of co-presence (Co-Presence: Linguistic and Visual (LV) = -.5, Linguistic-only (L)= .5), Partner identity (Partner: B = - .5 vs. A = .5) and their interaction, using the maximal random effect structure possible. For each fixed effect and interaction, we report the unstandardized coefficient and its standard error, along with the associated z- or t-value and p-value, and the overall variance captured by the model (R²). Statistically significant predictors (at the Bonferroni adjusted $p = .0167$ level) are in bold. The variance and standard deviation of each random effect are reported in Table B2.

Table B2

The variance and standard deviation of each random effect of the models reported in Table B1 for total words (log-transformed), total idea units (log-transformed), the proportion of items with hedges, number of turns, and number of errors in Phase 1.

Total idea units				Total words		Items with hedges	
Random effects							
Groups	Name	Variance	SD	Variance	SD	Variance	SD
Triad	Intercept	0.0001	0.0087	.0006	0.0240	0.0000	0.0000
	Co-presence: L-only	0.1218	0.3490	.1127	0.3356	2.6420	1.6254
	Co-presence: LV	0.0854	0.2921	.0730	0.2702	0.6765	0.8225
	Partner A	0.0522	0.2284	.0688	0.2623	—	—
	Partner B	0.0000	0.0065	.0000	0.0045	—	—
Item	Intercept	0.0066	0.0815	.0000	0.0000	0.3973	0.6303
	Co-presence: L-only	0.0146	0.1207	.0135	0.1163	—	—
	Co-presence: LV	0.0363	0.1906	.0775	0.2783	—	—

		Number of turns		Number of errors	
		Variance	SD	Variance	SD
Triad	Intercept	0.0000	0.0016	.0000	0.0000
	Co-presence: L-only	9.4410	3.0726	.0199	0.1410
	Co-presence: LV	1.3990	1.1829	.0000	0.0054
	Partner A	2.1640	1.4712	.0072	0.0850
	Partner B	0.6380	0.7987	.0019	0.0439
Item	Intercept	0.0001	0.0118	.0000	0.0000
	Co-presence: L-only	1.0930	1.0454	.0125	0.1119
	Co-presence: LV	0.8785	0.9373	.0003	0.0180
	Partner A	1.7760	1.3325	.0023	0.0483
	Partner B	0.0248	0.1574	.0001	0.0079

Note. All models are based on 511 observations across 32 triads and 16 items (with half the items described to partner A and half to partner B, in each triad).

Appendix C

Table C1

The variance and standard deviation of each random effect of the models reported, in Table 4, for total idea units (log-transformed), total words (log-transformed), the number of reconceptualizations, and, in Table 5, for the presence of hedges, the presence of definiteness (definite expressions and meta-references), and the directors' description onset times for their initial descriptions in Phase 2.

		Initial idea units		Initial words		Reconceptualizations	
Random effects							
Groups	Name	Variance	SD	Variance	SD	Variance	SD
Triad	Intercept	0.0000	0.0001	0.1119	0.3346	0.0000	0.0002
	Linguistic Co-presence: No	0.1052	0.3244	0.0000	0.0000	2.6980	1.6427
	Linguistic Co-presence: Yes	0.0796	0.2821	0.0020	0.0449	0.4577	0.6765
	Visual Co-presence: No	—	—	—	—	0.4756	0.6896
	Visual Co-presence: Yes	—	—	—	—	0.0314	0.1773
Item	Intercept	0.0394	0.1986	0.0159	0.1142	0.0000	0.0000
	Linguistic Co-presence: No	0.0155	0.1244	0.0328	0.1883	1.6970	1.3029
	Linguistic Co-presence: Yes	0.0001	0.0096	0.0185	0.1464	0.3754	0.6127
	Visual Co-presence: No	—	—	—	—	0.0089	0.0941
	Visual Co-presence: Yes	—	—	—	—	0.0381	0.1953
		Hedges		Definiteness		Description onset times	
		Variance	SD	Variance	SD	Variance	SD
Triad	Intercept	0.0000	0.0000	0.0000	0.0000	0.0713	0.2671
	Linguistic Co-presence: No	0.0332	0.1823	0.0000	0.0000	—	—
	Linguistic Co-presence: Yes	0.3386	0.5819	0.2636	0.5134	—	—
	Visual Co-presence: No	1.1120	1.0550	0.4837	0.6955	—	—
	Visual Co-presence: Yes	0.3843	0.6199	0.3231	0.5684	—	—
	Partner A	0.0349	0.1867	0.3705	0.6087	—	—
	Partner B	0.1600	0.4001	1.2740	1.1290	—	—
	Partner Order: ABAB	0.0000	0.0000	0.0480	0.2190	—	—
	Partner Order: ABBA	0.0000	0.0000	0.5991	0.7740	—	—
Item	Intercept	0.0000	0.0000	0.0000	0.0000	0.0197	0.1403
	Linguistic Co-presence: No	0.0363	0.1905	0.0014	0.0380	—	—
	Linguistic Co-presence: Yes	0.3318	0.5760	0.2691	0.5188	—	—
	Visual Co-presence: No	0.5116	0.7153	0.0124	0.1115	—	—
	Visual Co-presence: Yes	0.0021	0.0459	0.2346	0.4844	—	—
	Partner A	0.0000	0.0000	0.0000	0.0000	—	—
	Partner B	0.0000	0.0000	0.0000	0.0000	—	—
	Partner Order: ABAB	0.0231	0.1518	—	—	—	—
	Partner Order: ABBA	0.0056	0.0752	—	—	—	—

Note. All models, except for description onset times, are based on 1020 observations across 32 triads and 16 items. The description onset times model is based on 865 observations across 30 triads and 16 items.

Table C2

The variance and standard deviation of each random effect of the models reported, in Table 6, for the number of turns (log-transformed) and errors in Phase 2.

			Turns		Errors	
Random effects						
Groups	Name		Variance	SD	Variance	SD
Triad	Intercept		0.0000	0.0007	0.0000	0.0000
	Linguistic Co-presence: No		0.0819	0.2861	0.0001	0.0104
	Linguistic Co-presence: Yes		0.0666	0.2582	0.0005	0.0234
	Visual Co-presence: No		0.0000	0.0030	—	—
	Visual Co-presence: Yes		0.0000	0.0029	—	—
Item	Intercept		0.0000	0.0001	0.0000	0.0006
	Linguistic Co-presence: No		0.0318	0.1784	—	—
	Linguistic Co-presence: Yes		0.0006	0.0241	—	—

Note. Both models are based on 1020 observations across 32 triads and 16 items.