

Lexical distributional cues, but not situational cues, are readily used to learn abstract
locative verb-structure associations.

Katherine E. Twomey^{1,2}, Franklin Chang^{2,3} & Ben Ambridge^{2,3}

¹Lancaster University, UK; ²ESRC International Centre for Language and
Communicative Development (LuCiD); ³University of Liverpool, UK

Address correspondence to Franklin Chang, University of Liverpool, Department of
Psychological Sciences, Eleanor Rathbone Building, Bedford Street South, Liverpool,
L69 7ZA, UK.

Email: franklin.chang@liverpool.ac.uk. Telephone: +44-151-794-1110

Author note

This research was supported by Leverhulme Research Project Grant RPG-158 to Ambridge and Chang. The authors are members of the International Centre for Language and Communicative Development (LuCiD). The support of the Economic and Social Research Council [ES/L008955/1] is gratefully acknowledged. We would like to thank Amy Bidgood, Bethan Ward, Leah Spinoza, Megan McCann, Sophie Owen, and Stacey George for their help with data collection and coding.

Abstract

Children must learn the structural biases of locative verbs in order to avoid making overgeneralisation errors (e.g., **I filled water into the glass*). It is thought that they use linguistic and situational information to learn verb classes that encode structural biases. In addition to situational cues, we examined whether children and adults could use the lexical distribution of nouns in the post-verbal noun phrase to assign novel verbs to locative classes. In Experiment 1, children and adults used lexical distributional cues to assign verb classes, but were unable to use situational cues appropriately. In Experiment 2, adults generalised distributionally-learned classes to novel verb arguments, demonstrating that distributional information can cue abstract verb classes. Taken together, these studies show that human language learners can use a lexical distributional mechanism that is similar to that used by computational linguistic systems that use large unlabelled corpora to learn verb meaning.

Keywords: lexical distributional learning; language acquisition; syntax acquisition; verb semantics; verb classes

Lexical distributional cues, but not situational cues, are readily used to learn abstract locative verb-structure associations.

1 Introduction

Language acquisition is a complicated business. With little explicit teaching from adults, children rapidly learn words and grammatical structures. Critically, children must acquire language-specific links between verbs and structures (Levin & Rappaport Hovav, 2005); for example, in English *fill* can appear in *the woman filled the bucket with water*, but not **the woman filled water into the bucket*. At around five years of age, children sometimes make *overgeneralisation errors* such as **I'm going to cover a screen over me* (4;5; Bowerman, 1982) where a verb is paired with a structure that is not appropriate in the language that they are learning. Such errors show that children understand the verb's meaning and can produce the structure, but they have not yet learned the correct verb-structure link. Over time, however, children stop making these errors. This *retreat from overgeneralisation* occurs as children learn adult-like verb-structure links (Pinker, 1989).

The English locative alternation (e.g., *I sprayed water onto the wall*) involves events where a *theme* (e.g., water) moves to a *location* (e.g., wall) and the location is changed by the action (e.g., wall becomes wet). Locative events can be described with two structures, which differ as to whether the verb is followed by the location or the theme: the location-theme (*LT*) structure, as in *the woman sprayed **the wall** with paint*, and the theme-location (*TL*) structure, as in *the woman sprayed **paint** onto the wall*. Not all locative verbs can appear in both structures, however. Specifically, *LT-biased* verbs appear predominantly in the *LT* structure, for example *deluge*, *inundate* and *flood* (e.g., *I deluged the flowerbed with water* vs. **I deluged water onto the flowerbed*). *TL-biased* verbs such as *dribble*, *drip* and *pour* appear mainly in the *TL*

structure (e.g., *I dribbled water onto the flowerbed* vs. **I dribbled the flowerbed with water*). Finally, alternating verbs like *spray*, *load* and *pack* appear in both structures (e.g., *I sprayed water onto the flowerbed/I sprayed the flowerbed with water*).

Linguistic analyses explain these associations between verbs and structures in terms of verb classes: clusters of verbs with common semantic and syntactic properties (Levin & Rappaport Hovav, 2005; Levin, 1985). For example, verbs in the “cover-type” class (e.g., *deluge*, *inundate* and *flood*; Ambridge, Pine, & Rowland, 2012; Pinker, 1989) have the semantic property “a layer completely covers a surface”, which highlights the surface location over the theme argument. The greater salience of the location in these actions means they tend to be described with utterances that place the location earlier in sentences (e.g., *he flooded the floor with water*). If verbs incorporate this relative salience information into their meaning representation (Levin, 1993), then the structural preferences of the verb can be determined from the meaning (e.g., location-salient verbs tend to appear in LT structures).

One potential solution to the problem of learning verb-structure mappings would be for children to learn conservatively, memorising the verb-structure mappings in their input. This could be implemented with statistical learning mechanisms such as entrenchment (e.g., Braine & Brooks, 1995) and preemption (e.g., Goldberg, 1995). Here, the occurrence of a particular verb in grammatical constructions (e.g., *I dribbled water onto the flowerbed*) constitutes probabilistic evidence for the mappings in the input and against the grammaticality of unwitnessed combinations (e.g., **I dribbled the flowerbed with water*). Although these proposals enjoy some support, including in this particular domain (Ambridge et al., 2012; Bidgood, Ambridge, Pine & Rowland, 2014) and can help to explain the retreat from

overgeneralisation errors, they are not on their own sufficient, as they do not directly explain why children make errors in the first place.

An influential account of why children overgeneralise is that of Pinker (1989), who suggested that from the outset children possess innate *broad range* rules that link alternating structures which can be used to describe the same action. In the locative, the broad range rule connects two construals of a locative action – one in which the focus is the location's change of state and the other in which the focus is the manner of motion of the theme. When the location's change of state is highlighted (e.g., a wall becomes completely covered with paint), the LT structure is preferred since it places the location earlier in the sentence (e.g., *the girl sprayed the wall with paint*). When the manner of motion is highlighted (e.g., the paint moves in a distributed manner under pressure), the TL structure is preferred, since it places the theme earlier in the sentence (e.g., *the girl sprayed the paint onto the wall*). On Pinker's account, the semantic information in the scene (e.g., thematic roles) can be used to activate a broad range rule that allows children and adults to take a verb that has been heard only in one structure and use it with the other structure. This is desirable for many low-frequency alternating verbs which may only have been in a single structure in the input (e.g., *strew the flowerbed with seeds* -> *strew seeds onto the flowerbed*), but it can also lead to overgeneralisations if a verb is only acceptable in one structure (e.g., **I filled the water into the bath*).

Pinker (1989) explains the retreat from overgeneralisation through the acquisition of semantic verb classes. In this theory, children assign verbs to semantic verb classes which link to structures via *narrow range* rules and these rules allow children to retreat from the overgeneralisations licensed by the broad range rules. In particular, the salience and consistency of the components of an action across

different instances determines its verb class (Gropen, Pinker, Hollander & Goldberg, 1991a). For example, LT-biased *cover* is used to describe an action where the location changes state from being visible to being obscured (e.g., a blanket covers a bed). While the state change is salient and consistent across different cover actions, the movement of the theme can take place in various ways (e.g. the blanket can be dragged, thrown, dropped, etc.). Likewise, TL-biased *pour* describes a liquid moving in a continuous stream to the location (e.g., water flowing out of a hose), but the change of state of the location can be variable (e.g., the bucket can be partially or fully filled; the water could be poured onto the floor, etc.). A range of empirical evidence supports the idea that for both adults and children, verbs' syntactic behaviour is governed by these semantically constrained classes (Ambridge et al., 2012a; Bidgood et al., 2014; Brooks & Tomasello, 1999; Brooks, Tomasello, Dodson & Lewis, 1999; Gropen, Pinker, Hollander, Goldberg & Wilson, 1989; Gropen et al., 1991a, 1991b; Pinker, Lebeaux & Frost, 1987).

Pinker's account of verb class acquisition focuses on semantic information that can be extracted from the situations that verbs are heard in. Gropen et al. (1991a) provided evidence in support of this *situational* approach in a series of verb learning experiments in which they taught children and adults novel verbs (e.g., *look*, *this is keating*) alongside novel actions. Each action included either a salient location change of state (e.g., a colour change, Exp 2) or a salient theme manner (e.g., moving a matchbox in a zigzagging motion, Exp 1). After training with these novel verb/action pairs, participants were prompted at test to describe the same action using a full locative structure. Participants used more LT locatives after training scenes with a salient location component, and more TL locatives after training scenes with a salient manner component. However, although this study appears to show situational

effects on verb-structure learning, this is not the only possible account of Gropen et al.'s results, because their test actions were biased in the same way as their training items. For example, if in training participants saw the theme move towards the location in a zigzag motion with no change to the location, they saw the same event again at test. Participants' choice of structure could therefore have been determined by placing the salient argument (location/theme) earlier in the sentence; importantly, this could take place without reference to verb-specific semantics (this experiment provided verb-independent constructional meaning; Goldberg, 1995; Twomey, Chang & Ambridge, 2014). More generally, since most studies that show semantic effects on structural choice manipulate the test situation (Ambridge et al., 2012; Bidgood, et al., 2014; Brooks & Tomasello, 1999; Brooks et al., 1999; Gropen et al. 1989; Gropen et al. 1991a), it is not clear whether learners can recall situational information previously associated with a verb and use that information in later structural choices. In Experiment 1, we examine whether verb-specific situational training information can influence structural choices at a later test.

A potential problem for situational learning is that the relevant situational information may only rarely be present: speakers do not generally narrate events as they unfold. Instead learners may acquire a considerable amount of information regarding a verb's meaning from its linguistic context, as proposed under the *syntactic bootstrapping* hypothesis (Gleitman 1990; Fisher, Gertner, Scott, & Yuan, 2010). For example, Naigles (1990) demonstrated that children correctly associated sentences containing novel verbs with causative visual scenes based on the transitive syntactic frame in which the verbs were presented. Specifically, children mapped the transitive sentence *the duck is gorpung the bunny* to a scene in which a duck made a bunny squat by pushing on the bunny's head (i.e., a causative action). In contrast, children

associated a scene in which a duck and a bunny simultaneously made arm gestures (i.e., a non-causative action) with intransitive sentences such as *the duck and the bunny are gorging* (Naigles, 1990) or *the duck is gorging with the bunny* (Hirsh-Pasek & Golinkoff, 1999; Kidd, Bavin, & Rhodes 2001). The results of syntactic bootstrapping studies have been explained with a range of distinct mechanisms. One involves the number of arguments in a phrase; for example, two arguments would signal a causative meaning (Fisher, 1996). Another account is that learners use syntactic structures to establish elements of verb meaning; for example, the sequence of syntactic categories NP VERB NP might bias towards the causative (Fisher, Gleitman, & Gleitman, 1991; Gleitman, 1990). A third account is that the post-verbal noun may signal its thematic role; for example, patient nouns may indicate the causative (e.g., Rappaport Hovav & Levin, 2002). In addition to these syntactic mechanisms, it has been suggested that lexical mechanisms could provide cues to verb meaning. For example, Mintz (2003) showed that frequent lexical frames could be used to classify words into categories like VERB or NOUN. Importantly, because frames like *is_the* will only pick out transitive verbs (e.g. *the boy is pushing the girl*), these frequent frames could be useful cues to verb meaning. Finally another mechanism is offered in Scott and Fisher (2009), who found that the lexical distribution of animate/inanimate subject pronouns in training could influence verb class acquisition. Thus, in contrast to the non-linguistic information used by situational theories of verb learning, syntactic bootstrapping approaches suggest that syntactic frames, thematic roles, arguments, lexical frames, and lexical distributions in the linguistic signal could support verb learning.

These syntactic bootstrapping studies have suggested several mechanisms which children may exploit when learning verb meanings. These different accounts

can be tested by examining verb learning in the locative alternation, as the alternation itself rules out some mechanisms. Since the LT and TL structures have three arguments and a similar surface structure (i.e., NP VERB NP PP), it would be difficult to use the number of arguments, syntactic structures, or frequent frames to learn verb classes. In addition, both post-verbal arguments are inanimate, so unlike in Scott & Fisher (2009), animacy/pronoun distribution is not a clear cue to locative verbs' structural biases. A further challenge comes from the fact that locative verbs do not always occur in locative structures in the input. Twomey et al. (2014) examined all utterances containing any of the 140 locative verbs examined by Ambridge et al. (2012) from all UK corpora in the CHILDES database of child-directed speech (MacWhinney, 2000), and found that 78% of adults' locative verbs in their sample occurred not in full locative structure, but in transitive or intransitive structures, for example *you dump the lady's toys*. Although the preposition is a very good cue for the locative structure (e.g., *with* for the LT structure; *into*, *onto* for the TL), this large corpus analysis showed that many locative verbs frequently did not occur with prepositions.

These features of the locative suggest that the acquisition of the structural properties of these verbs may depend heavily on lexical distributional learning. A growing literature in computational linguistics suggests that distributional learning mechanisms may provide a general account of lexical class learning (Brent, 1993; Cartwright & Brent, 1997; Dumais & Landauer, 1997; Merlo & Stevenson, 2001; Redington, Chater & Finch, 1998; Rohde, Gonnerman, & Plaut, 2009). For example, these algorithms could use the sentences *you've drenched **the carpet** with water* and *he saturated **his carpet*** to classify *drench* and *saturate* as being more similar to each other than they are to *fill*, which does not typically occur with *carpet*. Such models

achieve high levels of syntactic and semantic performance using the full set of words that occur with each verb (the lexical distribution). For example, Mikolov, Chen, Corrado & Dean (2013) demonstrated that lexical distributional regularities from six billion words from Google News were able to achieve state-of-the-art (at the time) performance in classifying pairs of words as being syntactically and semantically related. These mechanisms also work with child-directed speech: Scott and Fisher (2009) found that it was possible to distinguish causal and contact verbs using the distribution of subjects with these verbs in CHILDES corpora. In sum, there is growing evidence that distributional learning can be used to learn a range of meaning and syntactic distinctions from corpora.

Twomey et al. (2014) applied a distributional learning technique to the acquisition of locative verb classes. Their correspondence analysis (CA) used the words that appeared near a verb to classify it along several dimensions that encoded lexical distributional similarity. These dimensions predicted adults' verb grammaticality ratings (Ambridge et al., 2012; Bidgood et al., 2014). The CA was created from a list of the two post-verbal words for each verb in all of the parental input in the UK CHILDES corpora; the list could include any word, for example determiners, prepositions, nouns, verbs, and adjectives. The CA mapped these verbs into a similarity space based on the overlap in post-verbal words. Figure 1 provides an example of how this might work for a small set of verbs given a small set of nouns that might appear with them in corpora. In this example, the verbs *pour* and *inject* are close to each other in the similarity space, because they both occur with words like *water* and *oil* in the post-verbal position. The word *fill* is on the other side of the space, because it tends to have containers in post-verbal position. *Spray* is in the middle, because it alternates, and is sometimes followed by liquids, sometimes by

containers. *Load* is an LT-biased verb like *fill*, but since it has different nouns in post-verbal position, it is in a different part of the top part of the space. *Inoculate* is a low frequency verb that occurred only with one noun, *cow* (e.g., *they inoculated the cow*), but that is sufficient to place it close to *load*, which has also occurred with that noun. Notice that this can take place even if the person interprets *cow* as being the patient thematic role, rather than as a location thematic role. This illustrates how a CA takes a lexical distribution without thematic role information and places verbs in a similarity space such that regions of this space act like verb classes which can be associated with structures.

Twomey et al. (2014) tested a range of different CA learners on the locative verbs in their corpus. The best CA, which used two post-verbal words, explained 47% of the variance in the independent verb-structure ratings. In contrast, a CA that used all of the post-verbal words only explained 38% of the variance, because the order of the nouns was lost when all post-verbal words were collapsed together, blurring the distinction between LT and TL structures; for example, *the woman poured water into the tub* and *the woman filled the tub with water* would both have *water* and *tub* as post-verbal words. The success of the *two post-verbal words* CA suggests that a distributional learner should be sensitive to a small window of adjacent words in learning verb classes. Twomey et al. (2014) tested this prediction in a connectionist model that was biased for learning adjacent regularities. The model captured early overgeneralisation of locative verbs and the gradual retreat from overgeneralisation through the acquisition of locative verb classes. Critically, the model's input was designed so that it could only learn these locative verb classes from the two post-verbal nouns in transitive utterances. This corpus and modelling work demonstrated

that development of locative verb classes can be explained by a distributional learning mechanism combined with the transitive input that children hear.

While Twomey et al. (2014) provided support for lexical distributional learning of verb meaning with corpus- analyses and connectionist modelling, there is little experimental evidence that children can learn these classes in the same way as these models. One study, Scott and Fisher (2009), has shown that toddlers can use the lexical distributional information before a verb when it is the *only* cue for verb meaning, but it is less clear if such cues will drive verb class acquisition when they are post-verbal and when situational information is also present. In addition, toddlers may be limited in their ability to deal with experimental task demands, which might mean that they do not combine situational and distributional information consistently early in development. In the case of the locative, there is evidence that children learn verb classes later in development (e.g., Ambridge et al., 2012). Thus, exploring these mechanisms in older children removes some of these limitations by testing learning mechanisms at ages where the ability to learn from situational and distributional cues should be robust.

In summary, children need to learn semantic verb constraints on their structural choices. The acquisition of locative verb classes allows us to contrast a situational account of this process (Pinker, 1989) with a lexical distributional account (Scott & Fisher, 2009; Twomey et al., 2014), while controlling for other sources of information. In Experiment 1, we pit these two accounts against each other in children and adults and then further explore the properties of a lexical distributional mechanism in Experiment 2.

2 Experiment 1: Situational and Lexical Cues in Verb-Structure Linking

In the first study we taught participants novel locative verbs alongside animations of novel actions in a training session and then examined how children and adults would use these verbs at test. For example, in the training action depicted in Figure 2, a robot fills up its arms with an oil theme from the cylinder on the left and then goes towards the cone location on the right. It then shoots the oil in large balls towards the cone, filling it with the oil.

To examine the role of situational information in locative verb class acquisition, we manipulated the salience of the locations and themes in training. In the *location-salient* condition, the action in Figure 2 involved a large change in the location object and little change to the theme (e.g., Figure 2 shows the cone getting completely filled by balls of oil that are shot at close range). In contrast, in the *theme-salient* condition the motion of the theme was highlighted while the location was less changed (e.g. the balls of oil bounced on the floor on the way to the cone but the cone was only partially filled). If situational information is used to learn locative verb classes, then participants should remember whether the location or theme was salient for each novel verb and then use this at test to bias for the appropriate structure (e.g., location-salient verbs should appear in the LT structure).

To examine whether learners can use lexical distributional regularities to acquire verb classes, we described the training scenes with sentences that varied in whether the post-verbal noun was the location (L-transitive, e.g., *the robot was pabbing the cone*) or the theme (T-transitive, e.g. *the robot was pabbing the oil*). We used transitive frames because corpus work has suggested these structures are the main context for learning about locative verbs in the input to children (Twomey et al., 2014). We were interested in whether participants could use lexical distributional regularities in training to assign novel verbs to appropriate verb classes. For example,

if *pabbing* occurred with the post-verbal theme-like nouns *oil* and *water* (T-transitive), could participants use this information to select a TL structure with this verb at test?

After hearing all of the training scenes for each of the four actions, participants were shown the same four actions with novel pairings of objects. They saw these new videos and heard them described with an intransitive structure that mentioned the target novel verb (e.g., *the robot was pabbing*) and were encouraged to describe the scene. Gropen et al. (1991a) used test stimuli in which either the location or the theme was salient. Participants could therefore use this situational salience to select a structure and insert the novel verb after the structure had been planned. To force participants to use their memory of verb-specific situational regularities that were experienced earlier in training, our test scenes combined the salient version of motion of the theme with the salient version of the change in the location. For example, Figure 3 depicts the test item for *pabbing*, in which the robot shoots balls that bounce on the floor (theme-salient manner) and the cone fills completely (location-salient endstate). The test action was shown with two themes (e.g., oil and water) and two locations (e.g., cone and box), which helped to bias participants to producing full locatives in order to disambiguate which themes and locations were involved in the action. Since the same test event was used regardless of the situational/lexical distributional condition in training, an effect of those variables at test would require participants to have retained some memory of the training situation.

Because test scenes had salient theme and location components (e.g., balls of oil bounce, cone fills completely), another way to show a situational effect would be to use the consistency of location or theme across training and test; for example the training action in Figure 2 and test action in Figure 3 both show the cone being filled

completely. Pinker (1989) claimed that verb classes were defined by consistent situational components across exemplars. For example, LT-biased verbs like *fill* tend to describe situations with consistent location changes, (i.e., objects being filled). The predictions of the training/test situational consistency account are the same as the predictions of the relative saliency account: in both cases, a location-salient action on a particular training trial should yield more LT utterances at test than an action presented in theme-salient training situation.

To examine how these mechanisms change over development, we tested three age groups. Bowerman (1982; 1988) reported overgeneralisation errors at around age 5 (5;0: *Can I fill some salt into the bear?*), demonstrating that children at this age can insert verbs productively into their learned locative structures. Furthermore, Ambridge et al. (2012) found that 5-year-olds were sensitive to semantic constraints on their use of locative structures, which suggests that they already have some semantic knowledge that could constrain verb classes. Thus, we tested this age group to examine how lexical and situational cues are used as they learn produce locative structures. Although the ability to learn novel verb-structure links is likely to increase over development, whether the ability to use lexical or situational cues also changes over time is unclear. Thus, to examine how cue use changes over development, we also tested 9-year-olds and adults.

2.1 Method

2.1.1 Participants

Adult participants were 48 native British English-speaking undergraduate students aged 18 to 22 years. Data from a further five participants were excluded due to experimenter error (2), equipment error (1) or because participants were non-native

speakers of English (2). Adults were recruited through a university participation scheme and received course credit for taking part. Child participants were 51 5-year-old children (24 girls) and 55 9-year-old children (23 girls). All children were British English-speaking and were recruited from local primary schools. Data from a further 11 participants were excluded due to equipment error (age 5: 6; age 9: 5). Parents' prior consent was obtained and children received stickers for participation.

2.1.2 Design

The study crossed age (age 5/age 9/adult), situational training (location-salient/theme-salient; within-subjects) and lexical training (L-transitive/T-transitive; within-subjects) in a 3 x 2 x 2 design. Our dependent measure was structure produced (LT locative or TL locative).

2.1.3 Situational stimuli

Visual scenes consisted of animations of four scenes depicting a robot performing novel actions on a set of items, computer animated in Processing (Reas & Fry, 2005). In each scene a robot caused a theme item to move towards a location item, resulting in a change of state in the location (Figure 4). Each novel action was a combination of a cause-motion action and state-change action (Levin, 1993), which created a verb that was felicitous in both LT and TL constructions. In action A, the robot threw a sheet-like theme which opened up in mid-air and covered the location. In action B, the robot filled a large location object by shooting or bouncing large balls of liquid into it. In action C, the robot raised a large object upwards by spraying a stream of small particles into it. Finally, in action D the robot decorated a large object with a smaller one after carrying the smaller object to the larger object either with static arms or in an up/down pumping motion.

Training stimuli are depicted in Figure 4. As in Gropen et al. (1991a), we manipulated the manner and endstate components of each action. Our location/theme salience manipulations correspond to the situational elements that linguists argue are involved in the locative alternation (Levin & Rappaport Hovav, 2005; Beavers, 2010; Goldberg, 1995) and which have been manipulated in previous studies (Gropen et al 1991a). Specifically, each manner component consisted of either more or less motion of the theme. In actions A and C, the robot was positioned so that the theme moved either a short distance or a long distance (more-motion). In action B, the theme either went directly to the goal or bounced on the floor on the way to the goal (more-motion). In action D, the robot carried the theme at a consistent height or in an up-down zigzagging motion (more-motion). Each endstate component consisted of either more or less change to the location. In action A, the location was partially or completely covered (more-change). In action B, the box was partially or completely filled (more-change). In action C, the location was caused to levitate to either a low or a high level (more-change). In action D, the theme was embedded in the surface of the theme either deeply or on the surface (more-change). Thus, the theme-salient level of the situational manipulation combined the more-motion manner with the less-change endstate, while the location-salient level combined the more-change endstate and the less-motion manner. Since each novel action is not easily described by a single English locative verb that encoded both the state change and the manner of motion, Figure 4 provides a separate English gloss for the theme-salient and location-salient versions of each action. These glosses used locative verbs from Levin (1993) and were accepted as plausible descriptions of these actions by adult participants in the norming study described in Section 2.3.

Test stimuli are depicted in Figure 5, and consisted of the more-change and more-motion components of each action. For example, in the test scene for action A the location was completely covered (more-change), and the theme moved a long distance (more-motion). Critically, our stimuli were computer-generated and therefore the endstate and manner components used in test scenes were *identical* to the matching endstate and manner components used during training. Previous studies did not control the test stimuli in this way, because endstate and manner components of human actions (Gropen et al., 1991a) or hand-animated videos (Ambridge et al., 2012) are variable and difficult to equate across trials. To reduce item-specific effects, we used novel pairings of objects at test. For example, oil is shot into the cone in training (Figure 2), but water is shot into the cone at test (Figure 3).

2.1.4 Lexical stimuli.

Lexical training stimuli consisted of transitive sentences spoken by the experimenter (Figure 6). Each sentence contained one of four novel words selected as plausible action labels for English speakers: *cringing*¹, *pabbing*, *veeming*, and *zopping*. To investigate whether word co-occurrences could bias participants towards producing LT or TL locatives, training sentences occurred either with a post-verbal location noun or a post-verbal theme noun. *L-transitive* stimuli were sentences with a post-verbal noun that labelled the onscreen location-like object (e.g., *the robot was cringing the box*). *T-transitive* stimuli were sentences with a post-verbal noun that labelled the theme-like object (e.g., *the robot was cringing the water*; post-verbal nouns are provided in Figure 6). Participants heard each training sentence twice. At test, to ensure participants were not biased to produce a particular structure, they

¹ Pronounced as in *bringing*

heard the novel verb in an intransitive structure (e.g., *the robot was cringing*) and were encouraged to describe the scene with an utterance that mentioned the objects that were involved in the action.

In order to see any effect of situational or lexical training on LT/TL structural choice at test, we needed to induce participants to produce full locative structures. Thus, participants received two *warm-up* scenes at the start of the experiment and two *locative prime* trials before the test trials (see top two rows of Figure 6). These were designed to increase the production of full locatives equally across all situational and lexical training conditions. These trials depicted scenes involving a loading action and were described with both the LT structure (e.g., *the robot was loading the table with the box*) and the TL structure (e.g., *the robot was loading the box onto the table*). To balance order effects, we created eight counterbalanced lists, each consisting of two warm-up trials, followed by eight training trials, two locative prime trials, and finally eight test trials (20 trials in total). Situational/lexical training pairings (e.g., location-salient + L-transitive; location-salient + T-transitive) were counterbalanced across participants and verbs, as was LT/TL order of warm-up and locative prime sentences. Order of presentation of verbs was rotated between lists, and training and test trials in a given list were presented in the same order. Left-right position of objects was randomised across participants. On test trials, front-back position and first theme item used were counterbalanced across participants. Figure 6 depicts an example counterbalance list.

2.1.5 Procedure (Adult)

Participants were told that they would be shown a video of a robot on a spaceship who would be carrying out known and novel actions with items on the spaceship, and that their task was to describe the scene to the experimenter. The

experiment began with the two warm-up trials. First, the experimenter labelled the objects on the screen (location object first). Then, she said “Here is the robot in the spaceship. We’re going to watch him do some loading” and played the animation for the first warm-up trial. After the animation had finished playing, the experimenter presented the verb in one locative structure and asked an elicitation question (e.g., “The robot was loading the table with the box. The robot was loading the table with the box. Can you tell me what the robot was doing with the two things?”). The second warm-up trial depicted a second loading action using different objects, which was labelled in the same way using the alternate structure (e.g., “The robot was loading the pyramid onto the cart. The robot was loading the pyramid onto the cart. Can you tell me what the robot was doing with the two things?”). When participants produced a sentence which included only one of the objects, they were prompted to mention both objects up to three times, using the questions “Can you tell me what the robot did with both of the objects?”, “Can you tell me [*how/where*] the robot was [*verbing*] the [*object*]?” or using *onto/with*, “The robot was [*verbing*] the [*object*] [*onto/with*] what?” If the participant did not mention both objects after prompting, the experimenter noted this and started the next trial.

Eight training trials immediately followed the warm-up trials and proceeded in an identical manner except that the novel verbs were presented in the appropriate transitive sentence for the lexical condition in each counterbalance list (e.g., L-transitive: *the robot was veeming the ball*). Two locative prime trials followed the training trials in an identical manner to the warm-up trials. Finally, eight test trials were presented, again in the same way as training trials, with the exception that novel verbs were presented in intransitive form, for example *the robot was veeming*. When

participants mentioned only one of the objects, they were prompted in the same way as during training trials.

2.1.6 Procedure (Children)

The child procedure was the same as the adult procedure, with the following adaptations to ensure it was child-appropriate. First, in order to engage children in the task, before the experiment began the experimenter told the child she was looking for space scientists to help her with a special job on the spaceship. She then explained that she was preparing a report for the captain of a spaceship, that she would show the child a video of a robot doing special jobs using objects in the spaceship, and that the child's job was to watch very carefully and, when asked, tell her what the robot was doing so that she could complete her report. Second, warm-up, training and locative prime trials began with the experimenter labelling the objects on the screen and asking the child to repeat the labels, while on test trials the experimenter asked the children to label each object, correcting incorrect responses. In object labelling, the location object was always first in order to counter the strong TL bias found in children (Twomey et al., 2014). Since this object ordering was the same for all test items, it could only increase the overall use of LT structures, but cannot explain any variation due to situational or distributional manipulations in training. Third, to encourage children to use the novel verb, the experimenter repeated it in the elicitation question (e.g., "Can you tell me, using both things, what the robot was cringing?"). There were no other differences between the adult and child procedures.

2.2 Coding

Participants' responses on test trials were transcribed and coded offline. Sentences in which both nouns were used unambiguously were coded as *LT* or *TL* locatives (e.g., TL: *the robot was veeming the net onto the box*). 245 non-locative

sentences were coded as *Other* and excluded from further analyses (see Table 1; note that these responses were equally distributed across the combinations of situational and lexical distributional conditions, $\chi^2(1) = 0.26, p = .61$). Finally, data from 21 individual responses were excluded due to experimenter error on one or both training or test trials for that verb, equipment error or interruption (age 5: 11; age 9: 10; adults: 0). 959 locative responses were included in the final analysis (age 5: 210; age 9: 378; adult: 371). A further 25% of responses were coded by a second experimenter, naïve to the experimental hypotheses. Inter-coder reliability was substantial, Cohen's kappa = 0.92). The *how/where* and *onto/with* prompts were used to encourage full locatives. The *how/where* prompt was used on 2% (10) of the test trials with adults. No adults received the final *onto/with* prompt at test. The *how/where* prompt was used on 9% (18) of the test trials for 5-year-old children and 0.3% (1) of the test trials with a 9-year-old child. The *onto/with* prompt was used on 3% (6) of the test trials for 5-year-old children and 0.3% (1) of the test trials with a 9-year-old child. Although these prompts could bias towards particular structures, they were only used when the participant had already mentioned either the theme or location in the post-verbal position, indicating that that element was salient for the participant.

Table 1. *Non-locative responses excluded from analyses in Experiment 1.*

Type	Example	Age 5	Age 9	Adults
Intransitive	<i>the robot was pabbing</i>	10	1	0
L-transitive	<i>the robot was zopping the box</i>	16	2	0
T-transitive	<i>the robot was veeming the net</i>	23	9	4
Ambiguous transitive	<i>the robot was zopping the all of them</i>	2	0	0
LT with incorrect	<i>he was pabbing the cone onto the oil</i>	2	0	0

preposition				
TL with incorrect preposition	<i>he was zopping the oil with the box</i>	8	17	9
Other	<i>getting oil on hands and shooting; the net onto the box</i>	126	23	0

2.3 Results

Proportions of LT locatives out of all LT and TL utterances produced by 5- and 9-year-old children and adults after situational training and lexical training are presented in Figure 7. Structure produced (LT = 1, TL = 0) was submitted to a binomial mixed effects model (lme4 version 1.1-5; Bates, Maechler, & Walker, 2014) with age (age 5 = -1, age 9 = 0, adult = 1), situational training (location-/theme-salient; effect coded) and lexical training (L-transitive/T-transitive; effect coded) as crossed fixed effects and participant and verb as random effects. The maximal model that converged included random intercepts for participant and verb, with no random slopes (Barr, Levy, Scheepers, & Tily, 2013). All mixed effects models reported used this random effects structure unless stated otherwise. All analyses were repeated excluding responses produced after prompting, however the results were similar to the full data set and we only report the full results here.

Overall, participants were biased towards the TL structure (*TL bias*; overall proportion LT = 0.29; negative logit intercept = -3.81, $SE = 0.86$, $z(949) = -4.44$, $p < .0001$), consistent with previous studies (Bidgood et al., 2014; Bowerman, 1982; Gropen et al., 1991a; Laffut & Davidse, 2002; Twomey et al., 2014). However, this early TL bias disappeared over development as LT production increased across age groups (main effect of age: $\beta = 2.13$, $SE = 0.65$, $\chi^2(1) = 12.00$, $p = .00053$).

Participants were sensitive to the lexical distribution, producing more LT locatives after L-transitives than after T-transitives (main effect of lexical training: $\beta = 1.45$, $SE = 0.33$, $\chi^2(1) = 45.32$, $p < .0001$) and relying more heavily on this information over time (age by lexical training interaction: $\beta = 1.12$, $SE = 0.37$, $\chi^2(1) = 9.72$, $p = .0018$). Situational information affected participants' LT production overall; however in contrast to Gropen et al. (1991a), more LT locatives were produced after theme-salient training than after location-salient training (main effect of situational training: $\beta = -0.59$, $SE = 0.30$, $\chi^2(1) = 6.69$, $p = .0097$). Finally, LT production was marginally affected by different combinations of situational and lexical information (lexical training by situational training interaction: $\beta = 1.27$, $SE = 0.61$, $\chi^2(1) = 2.81$, $p = .094$). No other interactions reached significance (age by situational training by lexical training interaction: $\beta = -1.06$, $SE = 0.73$, $\chi^2(1) = 2.15$, $p = .14$; age by situation training interaction: $\beta = -0.54$, $SE = 0.36$, $\chi^2(1) = 1.39$, $p = .24$).

To investigate how development affected particular age groups, we applied a binomial mixed effects model to each age group with situational training and lexical training as crossed fixed effects to address three questions: 1) Which age groups would show a TL bias? 2) How would the effect of lexical training change over age? 3) Was the mismatching situational effect, which did not interact with age, carried mainly by adults? With respect to the TL bias, the omnibus analysis found a main effect for age, where LT production rose as participants got older. Our separate models found a TL bias in both 5-year-olds (proportion LT = 0.17; negative logit intercept $\beta = -4.51$, $SE = 2.21$, $z(204) = -2.04$, $p = .041$) and nine-year-olds (proportion LT = 0.22; negative logit intercept: $\beta = -10.58$, $SE = 1.60$, $z(361) = -6.60$, $p < .0001$), but not in adults (proportion LT = 0.43, logit intercept = -1.00 , $SE = 0.79$, $z(365) = -1.27$, $p = .21$). Thus, the test stimuli appear to be unbiased in adults,

but children prefer to describe them with TL structures (as in corpus studies; Twomey et al, 2014). The second question was what the significant interaction between age and lexical training says about lexical distributional learning at each age group. When we looked at the effect of lexical training in the separate models, we found that 5-year-olds showed no effect ($\beta = 0.11$, $SE = 0.73$, $\chi^2(1) = 0.001$, $p = 0.97$), 9-year-olds did show an effect ($\beta = 1.86$, $SE = 0.66$, $\chi^2(1) = 9.34$, $p = .0022$), and adults showed a large effect ($\beta = 2.51$, $SE = 0.46$, $\chi^2(1) = 44.79$, $p < .0001$). This suggests that the ability to use the lexical distribution to learn about verbs grows over development, although it was not sufficiently robust for the 5-year-old children to learn from a few trials. Our final question related to the mismatching situational effect, by which participants preferred LT structures for theme-salient events. There was no situation effect for the 5-year olds ($\beta = -0.71$, $SE = 0.68$, $\chi^2(1) = 1.30$, $p = 0.25$) or 9-year-olds ($\beta = 0.77$, $SE = 0.64$, $\chi^2(1) = 0.75$, $p = 0.38$). The adults, on the other hand, showed a robust mismatch effect ($\beta = -1.33$, $SE = 0.39$, $\chi^2(1) = 9.76$, $p = .0018$). Although there is no situation training by age interaction, the fact that this effect appears to be strongest in the adults suggests that it could be the result of some experiment-specific strategy that only appeared in adults and some children.

As the situational effect in adults was opposite to the one predicted, we carried out an additional norming study to insure that our training stimuli were biased in the predicted direction, where location-salient stimuli were best described by LT structures. In this norming study, twelve new participants saw one location-salient scene and one theme-salient scene for each novel action from the training stimuli in the above study. For each scene, participants were asked to choose between a description with an LT structure and another with a TL structure containing the verbs in the English glosses in Figure 4 (nouns were adapted for the particular stimuli). Four

counterbalancing lists varied the order of action, situation bias, and object pairings. The two scenes for each action included different objects and were separated by three trials. Importantly, the location-salient and theme-salient scenes for each action depicted similar motion and endstates, so the verbs were compatible with both versions, as illustrated in the results in Figure 8. For example, both *veem* actions were compatible with *cover* and *throw*, since the theme was thrown and the location covered in either variant; participants could not therefore have used the verbs alone to discriminate between the scenes. Further, although individually these verbs could bias participants towards either LT or the TL structure, the fact that one LT and one TL verb was present on every scene means that any potential distributional bias was equated. Thus, to distinguish situational bias in their choices, participants had to be sensitive to situational information, and specifically, gradations in manner of motion or endstate. A binomial mixed effects model was applied to LT production with centred situation bias. Participants and verb were entered as random effects with situation bias as a random slope for both. There was a main effect of situation bias ($\beta = 1.67$, $SE = 0.65$, $\chi^2(1) = 2.58$, $p < 0.01$). This shows that adults could indeed distinguish the location- and theme-salient stimuli across the verbs and that they preferred to label the location-salient scene with an LT structure and the theme-salient scene with the TL structure. The difference in LT proportion between location- and theme-salient actions was 25%. Thus, in line with Gropen et al (1991a, Exp 3 found an 8% difference for alternating verbs), our norming study demonstrates that adults preferred LT structures for location-salient scenes more than for theme-salient scenes. Thus, the opposite effect of situation bias with the same animations in the main study must be due to the way that memory encodes the link between situational information and verbs.

Twomey et al. (2014) found a general TL bias in structural choices in their corpus analysis, where 66% of adult locatives and 87% of child locatives used the TL structure. They argued that this bias could trigger children to overgeneralise LT-biased locative verbs into the TL structure. Since verbs that appeared with location-salient/L-transitive training in Experiment 1 should be described at test with an LT locative, any TL locatives produced for these verbs can be thought of as a type of TL overgeneralisation. To examine this, we used TL production compared to chance (0.50) as an index of TL overgeneralisation for these items (i.e., above-chance TL production illustrates systematic TL overgeneralisation). As predicted, children at both ages produced significantly more TL locatives after location-salient/L-transitive training than expected by chance (two-tailed exact binomial tests, age 5: 44/55, $p < .0001$; age 9: 70/97, $p < .0001$). Adults' TL production did not differ from chance (46/93, $p = 1$). Thus, for verbs where both situational and lexical information should have cued the LT structure, children nonetheless systematically overgeneralised those verbs to TL locatives, and even adults did not prefer the LT structure. Based on the effect of children's TL bias on locative production, we predicted that this preference for post-verbal theme-like nouns would extend to their production of other structures. Indeed, as expected, children produced more T-transitives than L-transitives overall ($\chi^2(1) = 3.92$, $N = 50$, $p = .047$). The TL bias also appeared to affect the errors that they produced: they produced structures with post-verbal theme nouns followed by *with* (e.g. **the robot was zopping the oil with the box*) more often than structures with post-verbal location nouns followed by *into/onto* (e.g., **the robot was zopping the box into the oil*; $\chi^2(1) = 19.59$, $N = 27$, $p < .0001$). Taken together, these data support the claim that the majority of children's early locative overgeneralisation errors reflect a general TL bias in normal sentence production (Twomey et al., 2014).

In contrast to the view that situational cues are used in learning verb classes (Pinker, 1989), adults and children in this study did not base their choice of LT and TL structure on the salience of the location or theme in training or the shared endstate or manner consistency across training and test items. Instead, adults and 9-year-old children remembered lexical co-occurrence information in transitive training sentences and used full locative structures at test that reflected verbs' lexical bias in training. This ability grew over development, as indexed by the effect size measure Cohen's d (Cohen, 1988) which factors out the variance and sample size associated with each age group: 5-year-old children did not use lexical cues appropriately ($d = 0.06$), 9-year-old children showed a small effect of lexical training ($d = 0.10$), and adults showed a medium effect of lexical training ($d = 0.53$). This is consistent with Pinker's (1989) claim that verb classes develop slowly, and demonstrates the gradual nature of the retreat from overgeneralisation.

2.4 Discussion

Experiment 1 examined whether participants' choice between the LT and TL locative at test reflected some verb related information in training. To exhibit these effects, participants must produce locative responses at test. Our participants did so for these novel actions, producing 78% locatives overall. Participants also used non-locative conjoined transitives (e.g., *getting the water and zopping the cone*), which conform to our instructions to mention both location and theme and are consistent with the transitive structures used in training. Overall, however, only 60 responses were non-locatives that mentioned both location and theme, which is small compared to the 958 full locatives produced, suggesting that these novel actions were best described using the locative structure.

Previous studies of locative use in 5- to 9-year-old children (Gropen et al., 1991a; 1991b) have found that children can use situational information about location and theme salience to choose between LT and TL structures. Our norming study showed that adults showed significant matching preferences for our training stimuli (e.g., LT was more preferred for location-salient events). In our verb learning study, however, children and adults did not easily store situational information with particular verbs in a way that could help them to assign them to the appropriate LT- or TL-biased verb class. Participants who saw a salient location with a completely filled cone were not more likely to use an LT structure at test (by placing the cone earlier in the sentence) compared to those that saw a partially filled cone. These results support the view that the situational effects reported in Gropen et al (1991a) are the result of biases in the test stimuli that directly influence structural choices independently of the particular verb and its class (Goldberg, 1995).

If children cannot quickly learn from situational input how to assign a verb to its class, how do they learn the many low frequency locative verbs that exist in languages (e.g., *encrusted*, *dapple*)? Twomey et al. (2014) found that post-verbal lexical distributional regularities in transitive sentences could be useful in learning locative verb classes for these low frequency verbs. The present study demonstrated that 9-year-old children and adults, but not 5-year olds, could combine lexical distribution with verbs to bias structural choices. We also found that this ability grew with age, consistent with the idea that the use of lexical distributional cues might depend on previously learned verb classes (Twomey et al., 2014).

While this study provides evidence for the role of post-verbal nouns in verb-structure choices, it is not certain that these results involve abstract verb classes. This is because in order to maximise the chance that children would recall situational and

lexical cues, the objects seen on test trials were the same as those seen in training. Both types of cue could therefore be learned in an item-specific manner. For example, participants could have used situational cues to learn that *pabbing* involves oil specifically, rather than theme-like objects in general. However, the lack of a situational consistency effect in children suggests that even when test stimuli included the same items as in training, children's memory of the action carried out with those items during training was not sufficient to constrain structural choice at test. Similarly, the lexical training effect seen in 9-year-old children and adults could have benefitted from this item overlap. For example, if during training a participant heard *the robot was pabbing the oil* with a scene including the oil and the box, at test they would see an action involving the oil and the cone. If the participant remembered the post-verbal nouns paired with *pabbing* during training (i.e., *oil*, *water*), they could use those nouns to begin the locative produced at test, triggering a TL structure (e.g., *the robot was pabbing the oil into the cone*). Hence in Experiment 1, it remains possible that our effect of lexical distribution was due to our older participants having learned a trigram like *pabbing the oil* rather than having learned an abstract verb class which would allow *pabbing* to be followed by any theme-like object and then any location-like object. Therefore, Experiment 2 examines whether learners can use the lexical distribution to assign abstract verb classes that can be generalised to new nouns at test.

3 Experiment 2: Lexical Distribution in Abstract Verb Class Assignment

The results of Experiment 1 suggest that participants may have been relying on a distributional mechanism to learn verb biases (see *Introduction*). However, there are many distributional learning mechanisms that could be used to learn abstract verb

classes. Experiment 2 tests the predictions of Twomey et al.'s (2014) connectionist model, which assigned verbs to classes from the distribution of nouns that occurred after the verbs in transitive structures. For example in the model's input, a verb might occur in transitive structures with container objects 75% of the time and liquid objects 25% of time. Based on this distribution, the verb would be assigned to an LT-biased class, since these structures tend to have containers in the post-verbal position. Critically, the model assigned verbs to LT- or TL-biased classes having only ever encountered them in transitive structures (never the locative) and without any situational semantics (e.g., thematic role salience). We were interested whether humans could learn verb classes under input conditions similar to the model.

In Experiment 2, we taught adults four novel verbs (*dacked*, *keefed*, *pilked* and *tifed*) without visual input. We trained participants with these verbs in transitive sentences (e.g., *the grandmother tifed the juice*) and then tested them by presenting the verb with three nouns (*tifed* + *cup/coffee/salesperson*) and asking them to generate a sentence (e.g., *the salesperson tifed the coffee into the cup*). The main manipulation was the set of nouns that occurred with that verb. The L-biased condition included object nouns that were typical locations (e.g., *The cleaner pilked the floor*) and the T-biased condition included object nouns that were typical themes (e.g., *The babysitter tifed the lemonade*). To investigate whether this learning mechanism was statistical in nature, both of these conditions included one item with the opposite type of noun (e.g., the L-biased condition had four location nouns and one theme noun). If the verb learning mechanism was not statistical, but simply recorded whether a location or a theme noun had occurred with that verb, then there should be no difference between conditions, since both types of nouns occurred with each verb. However, if human learners are sensitive to the relative frequency, like the

connectionist model, then they should prefer the structure that was consistent with the most frequent type of nouns that were paired with that verb in training.

3.1 Method

3.1.1 Participants

20 monolingual, English-speaking adults from the university community (21 – 32 years) participated for course credit or as volunteers.

3.1.2 Design

The study manipulated lexical distribution bias of novel verbs in a within-subject design (T-biased/L-biased).

3.1.3 Stimuli

Four novel words used in previous child language studies served as novel verbs, all of which were presented in the past tense: *dacked*, *kefed*, *pilked* and *tifed*. Forty transitive sentence frames were created, half with the L-transitive (e.g., *the man _____ the box*) and half with the T-transitive (e.g., *the man _____ the water*). All frames used different agent, location and theme nouns (see Table 2). Location nouns were containers (e.g., *bag*, *cup*, *suitcase*) or surfaces (e.g., *counter*, *noticeboard*, *roof*). Theme nouns were liquids (e.g., *antifreeze*, *beer*, *milk*), plural nouns (e.g., *toys*, *cushions*, *slates*) and mass nouns (e.g., *confetti*, *gravel*, *litter*). Table 2 provides examples of training and test order and stimuli for the first 22 trials. Biased training sets of five items were created for each novel verb. *L-biased* sets included one novel verb in four L-transitives and one T-transitive (*pilked* in Table 2 is L-biased), and *T-biased* sets included one novel verb in four T-transitives and one L-transitive (*tifed* in Table 2 is T-biased). To allow participants to generate their own structures, we presented verbs with three nouns (henceforth *test triples*). Each triple included unique

nouns for agent, location and theme, none of which had appeared earlier. There were 5 test items for each of the four verbs (20 test triples).

Table 2. Example training and test procedure for novel verb *tifed*.

Trial type	Stimulus
Warm-up	<i>filled + glass/water/man</i>
Warm-up	<i>poured + tea/mug/woman</i>
Training	<i>The caterer tifed the honey (T-transitive)</i>
Known	<i>sprayed + bubbles/spectator/clown</i>
Training	<i>The waitress tifed the jar (L-transitive)</i>
Known	<i>loaded + pasta/dish/cook</i>
Training	<i>The grandmother tifed the juice (T-transitive)</i>
Known	<i>spilled + milk/bib/baby</i>
Training	<i>The plumber tifed the antifreeze (T-transitive)</i>
Known	<i>draped + table/cloth/housekeeper</i>
Training	<i>The babysitter tifed the lemonade (T-transitive)</i>
Known	<i>covered + deck/nets/fisherman</i>
Training	<i>The stallholder pilked the basket (L-transitive)</i>
Test	<i>tifed + mug/cordial/boy</i>
Training	<i>The customer pilked the counter (L-transitive)</i>
Test	<i>tifed + tank/oil/sailor</i>
Training	<i>The driver pilked the petrol (T-transitive)</i>
Test	<i>tifed + doctor/syringe/medicine</i>
Training	<i>The cleaner pilked the floor (L-transitive)</i>
Test	<i>tifed + cup/coffee/salesperson</i>
Training	<i>The researcher pilked the clipboard (L-transitive)</i>
Test	<i>tifed + glass/wine/waitress</i>

Transitive training sentences alternated with test trials, with the constraint that all training sentences for a given verb appeared before the corresponding test triples. Participants initially saw two warm-up items consisting of two triples with known locative verbs (*filled* and *poured*). The first block of five test triples appeared with known locative verbs to further encourage locative test responses (*covered, draped, loaded, spilled, sprayed*; Pinker, 1989). These known trials were interleaved with the five transitive training sentences for the first novel verb (e.g., *The babysitter tified the lemonade*). In the second block, the next novel verb appeared in transitive training sentences (e.g., *The driver pilked the petrol*), while the first novel verb (e.g., *tified*) appeared with five more test triples (see Table 2). This procedure continued until all four novel verbs had been trained and tested. The last block interleaved the final novel test triples with five known verbs in transitive frames (*encrusted, scattered, slopped, speckled, spread*). Blocks of L-biased training items alternated with T-biased training items (counterbalanced across participants).

Practice and filler stimuli were selected to have balanced structural biases: four known verbs were from a non-alternating LT-biased verb class (*fill, cover, encrust, speckle*; Pinker, 1989), four verbs from a non-alternating TL-biased verb class (*pour, spill, scatter, slop*), and four verbs from alternating locative verb classes (*spray, load, drape, spread*). Two counterbalance lists were created that varied which novel verb appeared in each section, such that each verb occurred in both L and T structures.

3.1.4 Procedure

Participants were tested individually in a quiet room. Sentences were presented and participant responses recorded using a program written in Processing v.2.0 (Reas & Fry, 2005). Training sentences were presented one word at a time in

the centre of the screen. To ensure that participants paid attention to the training sentences, presentation of each word was self-paced by pressing the Enter key. When the sentence was completed, the prompt *SPEAK* appeared on the screen. Participants then pressed the spacebar while repeating the sentence aloud.

Test trials displayed the three nouns from the test triple and a novel verb in a diamond shape on the screen. Subject and verb were randomly placed in the bottom and right locations, and location and theme were randomly placed in the top and left locations. This configuration made it harder for participants to use typical English reading patterns or to develop experiment-specific ordering strategies. Participants were asked to formulate a sentence with those words and to press the Enter key when ready. The prompt screen then appeared, and participants were again asked to press the spacebar as they said the sentence aloud. The test trials were designed to allow participants the freedom to generate their own sentence from the four words. Participants were told that they could add words like articles or prepositions in order to make their sentence grammatical.

3.2 Coding

Participants' responses on test trials were coded offline. Locative sentences with two unambiguous post-verbal nouns were coded as either *LT* or *TL*. For example, the utterance *the student dacked the noticeboard with posters* (with a post-verbal location-like noun) was scored as *LT*, while *the student dacked posters onto the noticeboard* (with a post-verbal theme-like noun) was scored as *TL*. Transitive sentences (e.g., *the salesperson tified a cup of coffee*) and other non-locative uses of the novel verb (e.g., *the keefed tourist put film in his camera*) were coded as *Other* and excluded from further analyses (48 transitives, 20 non-locative). 332 locative responses were included in the final analysis. All responses were coded by a second

coder naïve to the experimental hypothesis. Inter-coder reliability was substantial (Cohen's kappa = 0.92).

3.3 Results

Figure 9 (right panel) depicts the proportion of LT locatives produced for LT- or TL-biased novel verbs in Experiment 2. LT production (LT = 1, TL = 0) was submitted to a binomial mixed effects model with lexical training (centred) as a fixed effect and participant and verb as random effects. The maximal model that converged included random intercepts for participant and verb, with no random slopes (Barr et al., 2013). Overall, participants produced LT and TL locatives approximately equally frequently (overall proportion LT utterances = 0.46; logit intercept = -0.20, $SE = 0.39$, $z(328) = -0.52$, $p = 0.60$). LT production was higher for L-biased verbs than for T-biased verbs (beta = 0.97, $SE = 0.27$, $\chi^2(1) = 13.06$, $p = .00030$). Thus, because nouns encountered on test trials did not appear on training trials, participants in Experiment 2 assigned verbs to abstract classes based on the set of nouns that they appeared with during training and then at test, these classes biased their structural choices.

Despite the strong evidence for the importance of lexical distributional information from both studies presented here, it remains possible that the visual scene information provided in Experiment 1 may have affected participants' ability to use lexical cues: that is, it may be easier for speakers to use lexical cues when they have already extracted thematic role information from the visual scene. To examine this possibility, we compared the results from the adult data in Experiments 1 and 2. The left-hand panel of Figure 9 depicts adults' LT production in Experiment 1 based on lexical training, that is, collapsed across situational consistency. We submitted the adults' proportion LT production from both experiments to a binomial mixed effects model with training bias and experiment as fixed effects (effect coded). The maximal

model that converged included by-participant random intercepts and slopes for training bias (Barr et al., 2013). Adults were not TL-biased (overall proportion LT utterances = 0.44; logit intercept = -0.29, $SE = 0.65$, $z(696) = -0.45$, $p = .65$). Overall, adults' LT production was higher for L-biased verbs than for T-biased verbs (beta = 1.17, $SE = 0.51$, $\chi^2(1) = 19.22$, $p < .0001$). However, there was no effect of experiment (beta = -0.44, $SE = 0.80$, $\chi^2(1) = 0.12$, $p = .73$) or interaction between training bias and experiment (beta = 0.90, $SE = 0.69$, $\chi^2(1) = 1.72$, $p = .19$). Thus, the visual situational information (e.g., thematic roles, salience information) encountered in Experiment 1 did not substantially increase participants' use of lexical cues relative to Experiment 2, which included no situational information. Other differences between the studies (e.g., blocked training and test vs. interleaving training and test, number of nouns paired with each verb, overlap in nouns between training and test, number of locative fillers) did not strongly modulate the results. Thus, the simplest explanation for the effect of lexical distribution in both studies is that participants used the distribution of nouns in the training phase to assign verbs to abstract classes. The division of nouns into theme and location nouns was based on intuition (e.g., lemonade is more likely to be a theme than a location), but our participants shared these intuitions and when presented with three arguments like cup, coffee, and salesman, they were more likely to use the novel verb with a similar type of argument (e.g., the salesman tified the coffee into the cup). That is, they were able to generalise a novel verb to a structure that it had never been paired with based on a verb class that was shaped by the distribution of nouns in the absence of situational cues.

4 General Discussion

Theories of language acquisition often assume that verb knowledge is acquired by combining situational information from the world with abstract syntactic structures. Computational models do not always use these types of information to learn about verbs, because it can be difficult to identify the relevant aspects of scenes or accurately construct syntactic structures. Instead, some of these models make extensive use of distribution of words in sentences to identify aspects of word meaning and syntactic preferences (Mikolov, et al. 2013). Lexical distributional learning could provide a unified approach to explaining a range of different phenomena in language acquisition and adult processing (Twomey et al., 2014). For example, Chang, Bock, and Goldberg (2003) suggest that thematic role-based structural priming in the locative alternation could be due to differences in the lexical distribution of themes and locations. These mechanisms can also explain the behaviours seen in syntactic bootstrapping studies (e.g., Chang, Dell, & Bock, 2006), because these studies manipulate the lexical items around novel verbs (e.g., verbs followed by *the* are likely to be causative, verbs followed by *with* or occurring at the end of a sentence are more likely to be non-causative). While these mechanisms seem to be useful for learning about verbs, relatively little experimental work, other than Scott & Fisher (2009), isolates the role of lexical distribution from syntactic or situational variation. The present work addresses this gap using the locative alternation. It is learned relatively late in development, which means that it can be examined within a production task at an age where children should be able to learn and generalise verbs outside of the laboratory.

The acquisition of locative verb classes is a puzzle, because most locative verbs do not appear in the full locative structure in the input (e.g., statistical learning of the association between the verb and the transitive construction cannot fully

explain LT/TL biases in locative generalisation). And since the locative alternation controls factors such as syntactic frames, post-verbal noun animacy, and number of arguments, the syntactic bootstrapping mechanisms are not easily applicable here. One mechanism that can address this issue is lexical distributional learning, which Twomey et al (2014) showed could explain how these locative verb classes could be learned from post-verbal words. Their connectionist model attempted to predict the post-verbal words in locative utterances. When these predictions were incorrect, the prediction error was used to modify the internal verb representations which generated the prediction. Over time, the model developed distinct verb classes for predicting different types of post-verbal words. However, these classes by themselves did not create structural biases. Rather, these classes became associated with LT-only and TL-only classes learned from frequent verbs (e.g., *fill*, *pour*) that occurred with situational information (e.g. thematic roles). If verbs occurred predominantly with post-verbal words that also occurred with the LT-only class, they were classified as close to LT-only verbs in their structural preferences. Similarly, verbs were closer to the TL-only class if they occurred frequently with the same post-verbal words. Thus, situational information was used to learn verb-structure associations for frequent non-alternating verbs and this information was sufficient to associate structural biases with alternating verb classes using post-verbal words without situational information like thematic roles (Fig.1).

This mechanism can explain the lexical training effect in our studies. In Experiment 1, the 9-year-old children and adults encoded the nouns that occurred after the verb in training for four different novel verbs and used this information later at test. In Experiment 2, adults learned lexical distributional regularities and generalised this knowledge to new nouns at test. The 5-year-old children in

Experiment 1 did not show an effect of lexical training, suggesting that the ability to learn from the lexical distribution changed with age. This increase in ability mirrors Twomey et al.'s (2014) connectionist model, where the separation between verb classes increased slowly as the model learned. Figure 10 shows how this developmental process might work for the age groups in our study. Spatial distance in the figure encodes verb class similarity, and each verb is an exemplar of a verb class.

Given the early evidence for lexical distributional learning (Scott & Fisher, 2009), we assume that children at each age can use lexical distributional regularities to map novel verbs to existing locative verb classes. However, because 5-year-olds have a single cluster of locative verbs, it is difficult for them to show a distinction at test between verbs that have occurred in L-transitive and T-transitive sentences during training. Older children and adults know more verbs, and importantly these verbs are more semantically distinct (middle and right panel in Figure 10). This predicts that adults and older children will show stronger systematic structural choice at test than younger children. Support for the differentiation of verb classes over time is provided in Ambridge et al.'s (2012) rating study, in which the strength of verb-class related semantic predictors increased with age. Although our results come from older children, Twomey et al. (2014) suggested that these same distributional mechanisms could explain a range of early effects in development. For example, syntactic bootstrapping studies typically manipulate post-verbal words (e.g., novel verbs followed by articles or pronouns are more likely to be placed near the causative verbs in the semantic space). Distributional learning over these words would yield results that are similar to the predictions based on syntactic structure.

Another important feature of our data is the TL bias in children. Twomey et al. (2014) argued that there was a close relationship between verb classes and the TL bias, where the TL bias resulted from an early inability to distinguish locative verb classes. The results in Experiment 1 support this claim. Adults used LT and TL structures equally often, which would emerge from a broad spread of verb classes which they could attach to both LT- or TL-biased verbs based on lexical distribution (Figure 10, right panel). However, 5-year-olds and 9-year-olds showed a TL bias, which would emerge from a clustering of known verbs in the TL-biased part of the space (Figure 10, left and middle panels). When verbs are clustered tightly as in the 5-year old panel in Figure 10, an LT-only verb like *cover* could be placed into the TL-biased part of the verb space, explaining why children sometimes make overgeneralisation errors where LT-only verbs are placed into TL structures at this age (4;5 **I'm going to cover a screen over me*; Bowerman, 1982). The predictions of this lexical distributional account are different from approaches in which lexical item-based knowledge forms the basis for early usage (Tomasello, 2003). For example, *fill* appears frequently in LT and *pour* is frequent in TL in child directed speech (Twomey et al., 2014), so under an item-based approach, children should be accurate at both of these frequent pairings from the outset. However, Gropen et al. (1991b, Exp 1) found that young children (2;6-3;5) produced errors by placing the verb *fill* in the TL structure as often as they created the correct pairing, even though they correctly used the verb *pour* only in the TL structure. A lexical distributional approach can explain these findings: the first cluster is TL-biased, while *fill* is at the border between the TL and LT spaces (left panel of Figure 10).

Nonetheless, in addition to distributional information, there is substantial evidence that children understand aspects of situational meaning from early in

development (causality, Leslie & Keeble, 1987; goals, Luo, 2011) and that structural knowledge is linked to this semantic information (Ambridge, Pine, Rowland, Freudenthal & Chang, 2014; Ambridge et al., 2012a; Ambridge et al 2012b; Ambridge, Pine, Rowland, Jones & Clark, 2009; Bidgood et al., 2014). However there is relatively little experimental evidence that children can store structure-relevant situational features for novel verbs across separate training and test events (for an exception, see Ambridge, Pine & Rowland, 2011). Indeed, the effect of situation in the current study was not in the predicted direction: participants preferred to use LT structures with verbs that had been seen in theme-salient situations in training. Although there was no interaction with age, this effect was strongest in adults. In contrast, when adults were queried about the situational manipulation in the norming study, they produced structures which matched the situation, preferring LT structures for location-salient scenes. This suggests that adults can identify the salient elements in the training scenes and have a preference for the matching structure, but by the time they are tested, this preference has changed into a mismatch preference. The source of this mismatch effect is not clear, but one possibility is that the salience of the theme and location in the test situation is influenced by variability, rather than consistency as Gropen et al (1991a) argued. That is, endstate/manner in the situational manipulation varied between training and test, which could highlight the item that changed, and increase production of the structure that placed the highlighted item earlier. For example, the difference between a partially filled cone in training and a fully filled cone at test could render the cone more salient relative to the unchanged theme motion and trigger an increase in LT production at test. Regardless of what caused this mismatch effect, however, our results suggest that situation

information is not being transparently associated with verbs to create matching structural preferences.

An alternative mechanism that participants could have employed in these studies is to activate thematic roles directly from the lexical items and create verb classes based on these word-derived roles. For example, hearing the word *floor* in *the cleaner pilked the floor* could activate a location thematic role, and *pilked* could therefore be assigned to a location-biased verb class. Then at test, this verb class would bias speakers towards LT descriptions. One question is whether these word-derived thematic roles are linked to the situation-derived roles. If they are linked, the lexical effect should be stronger in Experiment 1, where additional visual information about thematic roles was provided, compared to Experiment 2, where no visual information was provided. However, there was no difference in the magnitude of the lexical effect between our experiments: word-derived roles plus situation-derived roles were not better than word-derived roles only. Linking word- and situation-derived thematic roles also predicts a strong interaction between situation and lexical training in Experiment 1. When the visual scene highlights the location role and the post-verbal noun activates the same location role, the location role should be very salient and easily bound to the verb. When the roles activated by the visual scene and the role activated by the post-verbal noun mismatch, then it should be harder to select the verb class, because the situation-derived and word-derived roles bias in opposite directions. This predicts that mismatching conditions should have a smaller lexical training difference than the matching conditions. In fact, the opposite was true. When situational salience matched the role selected by the lexical nouns, the difference between L- and T-transitives was 8%, but, when they mismatched, the lexical training difference was 18%. Thus, our results suggest that word-derived role

information is independent of situation-derived role information. This is consistent with the approach in Twomey et al.'s (2014) correspondence analysis, where the verb class space which encoded thematic role distinctions despite receiving only lexical – but not situational/visual - input.

Low frequency verbs are challenge for theories of language acquisition (e.g., *encrust, festoon, imbue, lard, mottle, replenish, shroud, wad*; Levin, 1993). The acquisition of the structural biases of these verbs is made more difficult when the verbs do not occur in the appropriate structures in the input, as is the case for many locative verbs. Lexical distributional learning can use overlap in adjacent words to identify semantically-related verbs, which can support the acquisition of these low frequency verbs, by using a small number of exemplars as in the studies presented here. In contrast, in Experiment 1 participants could not use situational information from a small number of exemplars to constrain verb meaning, even though the theme motion and location change were clearly visible and less variable than those in the real world (e.g., *the general imbued the soldiers with courage, the man infused the cake with vanilla*). Thus, while situational information should clearly be encoded with verbs when frequent or salient enough, the vast majority of linguistic forms are infrequent (Zipf, 1949). A complete theory of language acquisition requires a mechanism that can address the long tail of linguistic knowledge. The current studies point to the importance of lexical distributional learning in providing just such a mechanism.

References

- Ambridge, B., Pine, J. M., & Rowland, C. F. (2011). Children use verb semantics to retreat from overgeneralization errors: A novel verb grammaticality judgment study. *Cognitive Linguistics*, 22(2), 303–323. doi:10.1515/cogl.2011.012,
- Ambridge, B., Pine, J. M., & Rowland, C. F. (2012). Semantics versus statistics in the retreat from locative overgeneralization errors. *Cognition*, 123(2), 260–279. doi:10.1016/j.cognition.2012.01.002
- Ambridge, B., Pine, J. M., Rowland, C. F., Freudenthal, D., & Chang, F. (2014). Avoiding dative overgeneralisation errors: semantics, statistics or both? *Language, Cognition and Neuroscience*, 29(2), 218–243.
<http://doi.org/10.1080/01690965.2012.738300>
- Ambridge, B., Pine, J. M., Rowland, C. F., Jones, R. L., & Clark, V. (2009). A semantics-based approach to the “No Negative Evidence” Problem. *Cognitive Science*, 33(7), 1301–1316.
- 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. doi:10.1016/j.jml.2012.11.001
- Bates, D., Maechler, M., & Walker, S. (2014). lme4: Linear mixed-effects models using Eigen and S4. (Version R package version 1.1-5.). Retrieved from <http://CRAN.R-project.org/package=lme4>
- Beavers, J. (2010). The structure of lexical meaning: Why semantics really matters. *Language*, 86(4), 821–864.
- Bidgood, A., Ambridge, B., Pine, J. M., & Rowland, C. F. (2014). The retreat from locative overgeneralisation errors: A novel verb grammaticality judgment study. *PLoS ONE*, 9(5), e97634. doi:10.1371/journal.pone.0097634

- Bowerman, M. (1982). Reorganizational processes in lexical and syntactic development. *Language Acquisition: The State of the Art*, 319-346.
- Bowerman, M. (1988). The “no negative evidence” problem: How do children avoid constructing an overly general grammar. *Explaining Language Universals*, 73–101.
- Brent, M. R. (1993). From grammar to lexicon: unsupervised learning of lexical syntax. *Computational Linguistics*, 19(2), 243–262.
- Brooks, P. J., & Tomasello, M. (1999). How children constrain their argument structure constructions. *Language*, 75(4), 720–738.
- Brooks, P. J., Tomasello, M., Dodson, K., & Lewis, L. B. (1999). Young children’s overgeneralizations with fixed transitivity verbs. *Child Development*, 70(6), 1325–1337. doi: 10.1111/1467-8624.00097
- Cartwright, T. A., & Brent, M. R. (1997). Syntactic categorization in early language acquisition: Formalizing the role of distributional analysis. *Cognition*, 63(2), 121–170. doi:10.1016/S0010-0277(96)00793-7
- Chang, F., Bock, K., & Goldberg, A. E. (2003). Can thematic roles leave traces of their places? *Cognition*, 90(1), 29–49.
- Chang, F., Dell, G. S., & Bock, K. (2006). Becoming syntactic. *Psychological Review*, 113(2), 234–272.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Dumais, S. T., & Landauer, T. K. (1997). A solution to Plato’s problem: The latent semantic analysis theory of acquisition, induction and representation of knowledge. *Psychological Review*, 104(2), 211–240. doi: 10.1037/0033-295X.104.2.211

- Fisher, C. (1996). Structural limits on verb mapping: The role of analogy in children's interpretations of sentences. *Cognitive Psychology*, *31*(1), 41–81. doi:10.1006/cogp.1996.0012
- Fisher, C., Gertner, Y., Scott, R. M., & Yuan, S. (2010). Syntactic bootstrapping. *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*(2), 143–149. doi:10.1002/wcs.17
- Fisher, C., Gleitman, H., & Gleitman, L. (1991). On the semantic content of subcategorization frames. *Cognitive Psychology*, *23*(3), 331–392. doi: 10.1016/0010-0285(91)90013-E
- Gillette, J., Gleitman, H., Gleitman, L., & Lederer, A. (1999). Human simulations of vocabulary learning. *Cognition*, *73*(2), 135–176. doi: 0.1016/S0010-0277(99)00036-0
- Gleitman, L. (1990). The structural sources of verb meanings. *Language Acquisition*, *1*(1), 3–55. doi: 10.1207/s15327817la0101_2
- Goldberg, A. E. (1995). *Constructions: A construction grammar approach to argument structure*. Chicago: University of Chicago Press.
- Gropen, J., Pinker, S., Hollander, M., & Goldberg, R. (1991a). Affectedness and direct objects: The role of lexical semantics in the acquisition of verb argument structure. *Cognition*, *41*(1), 153–195. doi:10.1016/0010-0277(91)90035-3
- Gropen, J., Pinker, S., Hollander, M., & Goldberg, R. (1991b). Syntax and semantics in the acquisition of locative verbs. *Journal of Child Language*, *18*(1), 115–151. doi: 0.1017/S0305000900013325
- Hirsh-Pasek, K., & Golinkoff, R. M. (1999). *The origins of grammar: Evidence from early language comprehension*. Cambridge, MA: MIT Press

- Hohenstein, S., & Kliegl, R. (2013). Remef (REMove effects; version v0.6.10) [Software]. Available from <http://read.psych.uni-potsdam.de/joomla/attachments/article/12/remef.v0.6.10.R>
- Hovav, M. R., & Levin, B. (2002). Change of state verbs: implications for theories of argument projection. In Erteschik-Shir & Rapoport (Eds). *The Syntax of aspect: Deriving thematic and aspectual interpretation*. Oxford, UK: Oxford University Press. 274–286.
- Kidd, E., Bavin, E. L., & Rhodes, B. (2001). Two-year-olds' knowledge of verbs and argument structures. In M. Almgren, A. Barreñ a, M.-J. Ezeiz- abarrena, I. Idiazabal, & B. MacWhinney (Eds.), *Research on child language acquisition: Proceedings of the 8th conference of the International Association for the Study of Child Language* (pp. 1368–1382). Somerville, MA: Cascadilla Press.
- Laffut, A., & Davidse, K. (2002). English locative constructions: An exercise in neo-Firthian description and dialogue with other schools. *Functions of Language*, 9(2), 169–207. doi:10.1075/fo1.9.2.04laf
- Leslie, A. M., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition*, 25(3), 265–288. doi:10.1016/S0010-0277(87)80006-9
- Levin, B. (1985). Lexical semantics in review: An introduction. In Levin, B. (Ed.), *Lexical semantics in review: Lexicon project working papers #1* (Vol. 1). Cambridge, MA: MIT Center for Cognitive Science.
- Levin, B. (1993). *English Verb Classes and Alternations: A Preliminary Investigation*. London: University of Chicago Press.
- Levin, B., & Hovav, M. R. (2005). *Argument realization*. Cambridge: Cambridge University Press.

- Luo, Y. (2011). Three-month-old infants attribute goals to a non-human agent. *Developmental Science, 14*(2), 453–460. doi: 10.1111/j.1467-7687.2010.00995.x
- MacWhinney, B. (2000). *The CHILDES Project: Tools for Analyzing Talk, Volume II: The Database* (Vol. 2). Lawrence Erlbaum.
- Merlo, P., & Stevenson, S. (2001). Automatic verb classification based on statistical distributions of argument structure. *Computational Linguistics, 27*(3), 373–408. doi:10.1162/089120101317066122
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv Preprint arXiv:1301.3781*. Retrieved from <http://arxiv.org/abs/1301.3781>
- Mintz, T. H. (2003). Frequent frames as a cue for grammatical categories in child directed speech. *Cognition, 90*(1), 91–117. doi:10.1016/S0010-0277(03)00140-9
- Naigles, L. (1990). Children use syntax to learn verb meanings. *Journal of Child Language, 17*(02), 357–374. doi:10.1017/S0305000900013817
- Pinker, S., Lebeaux, D. S., & Frost, L. A. (1987). Productivity and constraints in the acquisition of the passive. *Cognition, 26*(3), 195–267. doi:10.1016/S0010-0277(87)80001-X
- Pinker, S. (1989). *Learnability and cognition: The acquisition of argument structure*. Cambridge, Mass.: Harvard University Press.
- Reas, C., & Fry, B. (2005). Processing.org: a networked context for learning computer programming. In *ACM SIGGRAPH 2005 Web program* (p. 14). Retrieved from <http://dl.acm.org/citation.cfm?id=1187351>

- Redington, M., Chater, N., & Finch, S. (1998). Distributional information: A powerful cue for acquiring syntactic categories. *Cognitive Science*, 22(4), 425–469. doi:10.1016/S0364-0213(99)80046-9
- Rohde, D. L., Gonnerman, L. M., & Plaut, D. C. (2005). An improved model of semantic similarity based on lexical co-occurrence. *Communications of the ACM*, 8, 627–633
- Scott, R. M., & Fisher, C. (2009). Two-year-olds use distributional cues to interpret transitivity-alternating verbs. *Language and Cognitive Processes*, 24(6), 777–803. doi:10.1080/01690960802573236
- Twomey, K. E., Chang, F., & Ambridge, B. (2014). Do as I say, not as I do: A lexical distributional account of English locative verb class acquisition. *Cognitive Psychology*, 73, 41–71. doi: 10.1016/j.cogpsych.2014.05.001
- Zipf, G. K. (1949). *Human Behavior and the Principle of Least Effort*. Boston, MA: Addison-Wesley.

Figure Captions

Figure 1: A correspondence analysis of the mapping between input distribution and verb semantic similarity

Figure 2. Example training action: robot fills a cone by shooting balls of oil at it

Figure 3. Example test action: robot fills cone with water by shooting balls that bounce on the floor

Figure 4. Location-salient/theme-salient training scenes for each novel action with situation-appropriate English gloss

Figure 5. Test scenes for each novel action

Figure 6. Example counterbalance list

Figure 7. Proportion of LT locatives produced for situational and lexical training in Experiment 1. Error bars represent standard errors after removing random effects from model (Hohenstein & Kliegl, 2013).

Figure 8. Proportion LT structures from corpora for theme and location-based nouns

Figure 9. Proportion LT production by lexical training bias for Exp. 1 and 2. For comparison, left panel depicts adult responses only in Exp. 1, pooled across situational consistency. Error bars represent standard errors after removing random effects from model (Hohenstein & Kliegl, 2013).

Figure 10. A developmental account of locative verb class acquisition

Figure1

[Click here to download high resolution image](#)

Input distribution

Verb	Post-verbal words
FILL	the, tub, cup, bottle
INOCULATE	cow
LOAD	the, wagon, cow
SPRAY	water, wall, tub
INJECT	oil, water, cow
POUR	the, water, oil, coffee

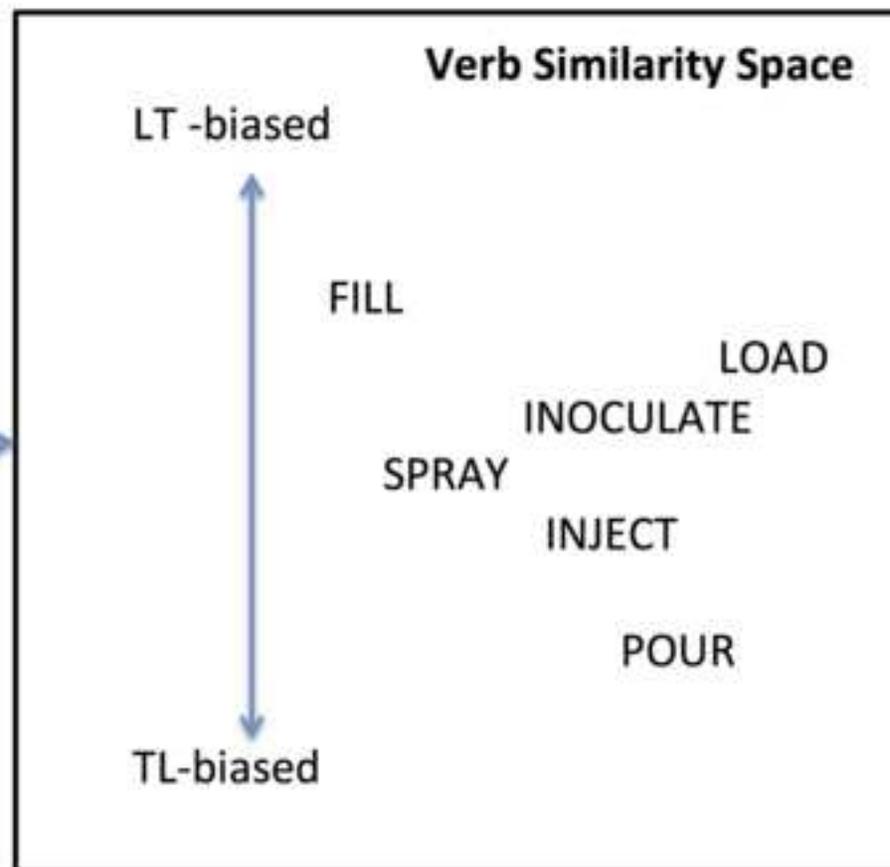


Figure2

[Click here to download high resolution image](#)

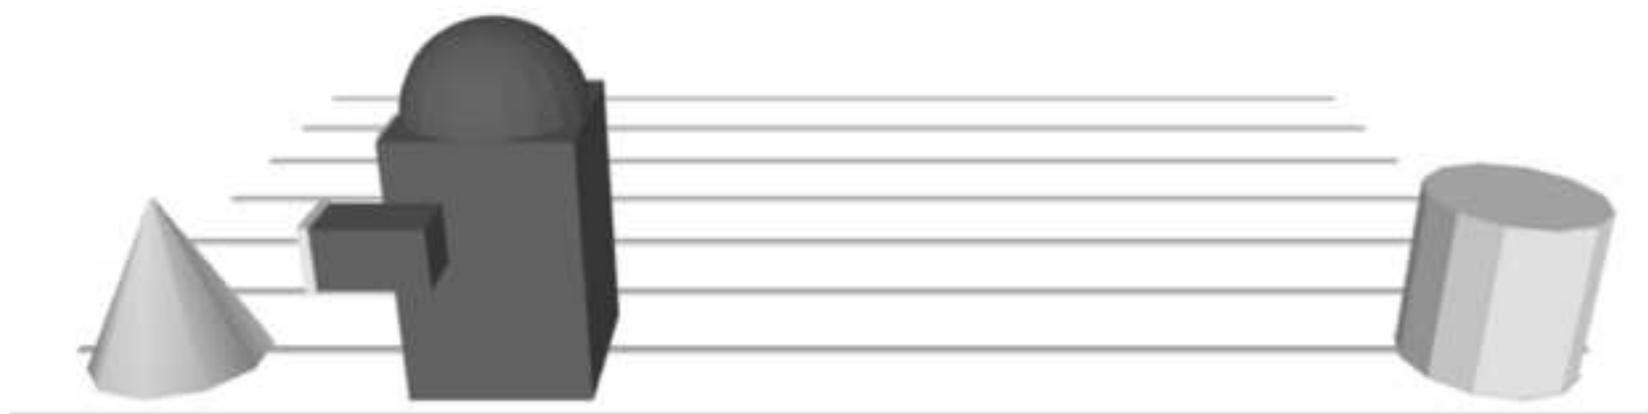
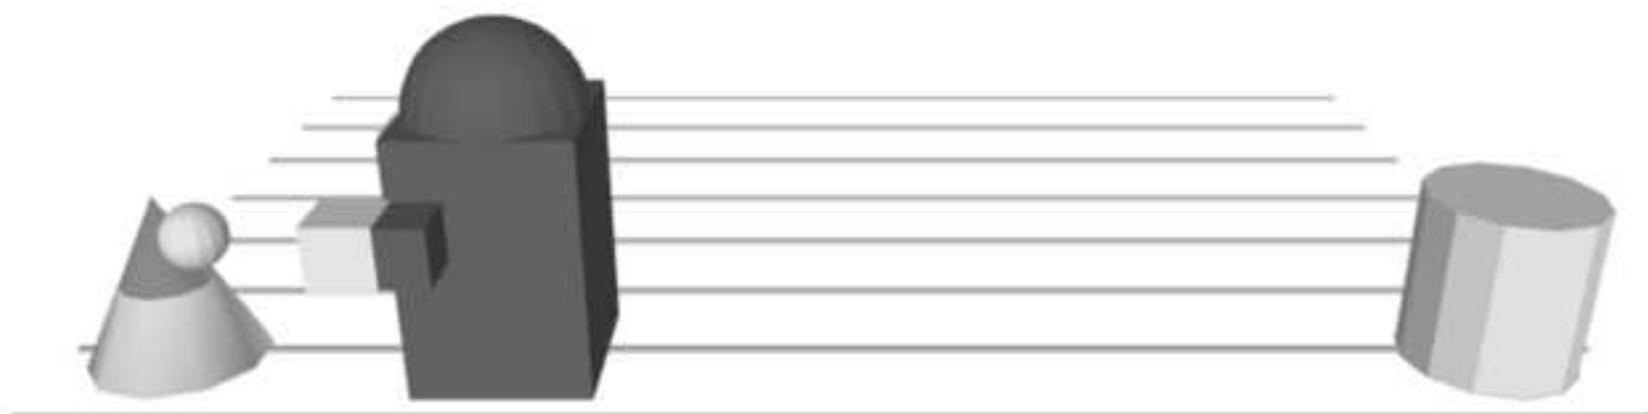


Figure3
[Click here to download high resolution image](#)

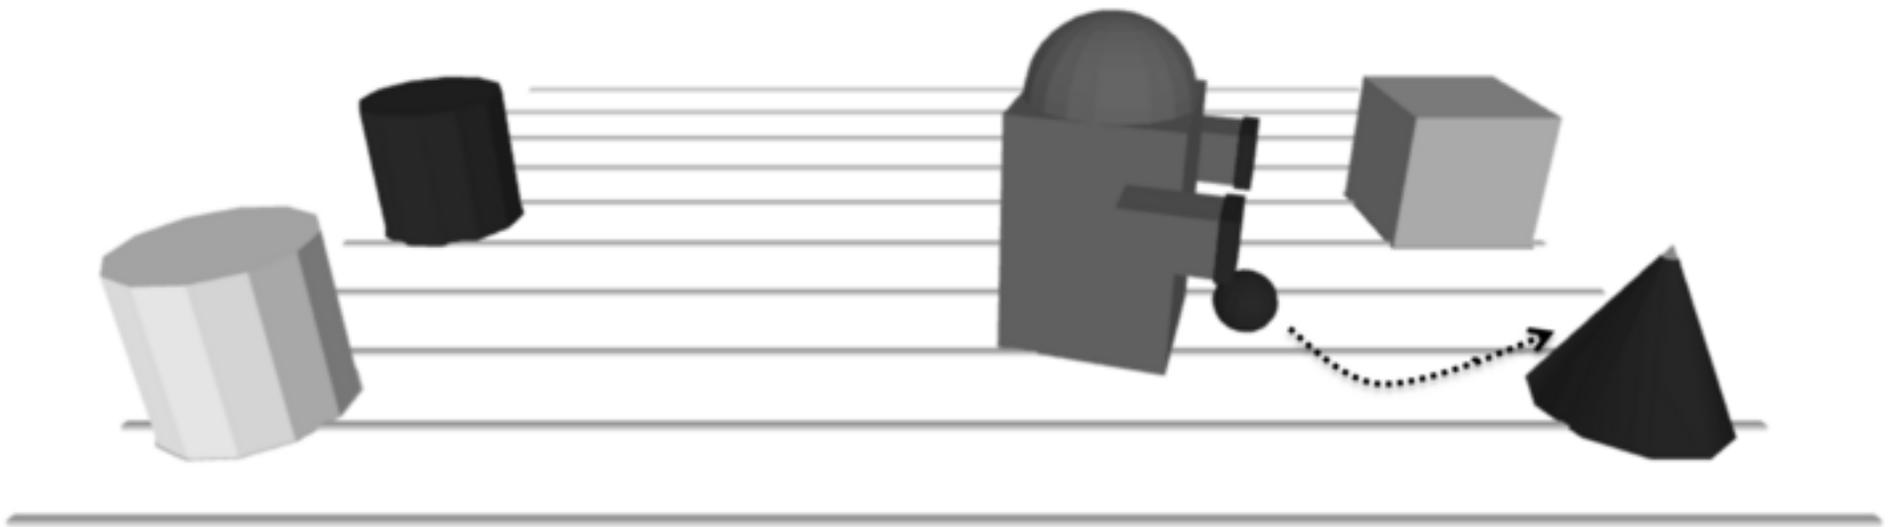


Figure 4

[Click here to download high resolution image](#)

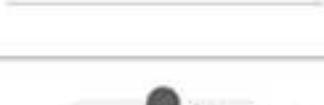
Action	Situational consistency	Middle of Event (theme motion)	End of Event (location change)
A	Location-salient Theme moves short distance Location completely covered <i>He covered the ball with the net</i>		
	Theme-salient Theme moves long distance Location partially covered <i>He threw the net over the ball</i>		
B	Location-salient Theme moves short distance Location completely filled <i>He filled the cone with the oil</i>		
	Theme-salient Theme bounces on floor Location partially filled <i>He shot the oil into the cone</i>		
C	Location-salient Theme moves short distance Location levitates to high level <i>He raised the box with the water</i>		
	Theme-salient Theme moves long distance Location levitates to low level <i>He sprayed the water on the box</i>		
D	Location-salient Theme moves in straight line Location conspicuously decorated <i>He decorated the box with the diamond</i>		
	Theme-salient Theme moves up and down Location subtly decorated <i>He embedded the diamond into the box</i>		

Figure5

[Click here to download high resolution image](#)

Action	Situational features	Middle of Event (theme motion)	End of event (location change)
A	Theme-salient: Theme moves long distance Location-salient: Location completely covered		
B	Theme-salient: Theme bounces on floor Location-salient: Location completely filled		
C	Theme-salient: Theme moves long distance Location-salient: Location levitates to high level		
D	Theme-salient: Theme moves up and down Location-salient: Location conspicuously decorated		

Figure 6

[Click here to download high resolution image](#)

Trial	Situational stimulus	Lexical stimulus
Warm-up 1		<i>The robot was loading the box onto the table</i>
Warm-up 2		<i>The robot was loading the cart with the pyramid</i>
Training 1 (location-salient)		<i>The robot was vacuuming the ball (L-transitive)</i>
Training 2 (location-salient)		<i>The robot was vacuuming the box (L-transitive)</i>
Training 3 (theme-salient)		<i>The robot was pushing the cone (L-transitive)</i>
Training 4 (theme-salient)		<i>The robot was pushing the box (L-transitive)</i>
Training 5 (location-salient)		<i>The robot was cringing the cross (T-transitive)</i>
Training 6 (location-salient)		<i>The robot was cringing the diamond (T-transitive)</i>
Training 7 (theme-salient)		<i>The robot was zipping the oil (T-transitive)</i>
Training 8 (theme-salient)		<i>The robot was zipping the water (T-transitive)</i>
Locative prime 1		<i>The robot was loading the table with the pyramid</i>
Locative prime 2		<i>The robot was loading the box onto the cart</i>
Test 1		<i>The robot was vacuuming</i>
Test 2		<i>The robot was vacuuming</i>
Test 3		<i>The robot was pushing</i>
Test 4		<i>The robot was pushing</i>
Test 5		<i>The robot was cringing</i>
Test 6		<i>The robot was cringing</i>
Test 7		<i>The robot was zipping</i>
Test 8		<i>The robot was zipping</i>

Figure 7

[Click here to download high resolution image](#)

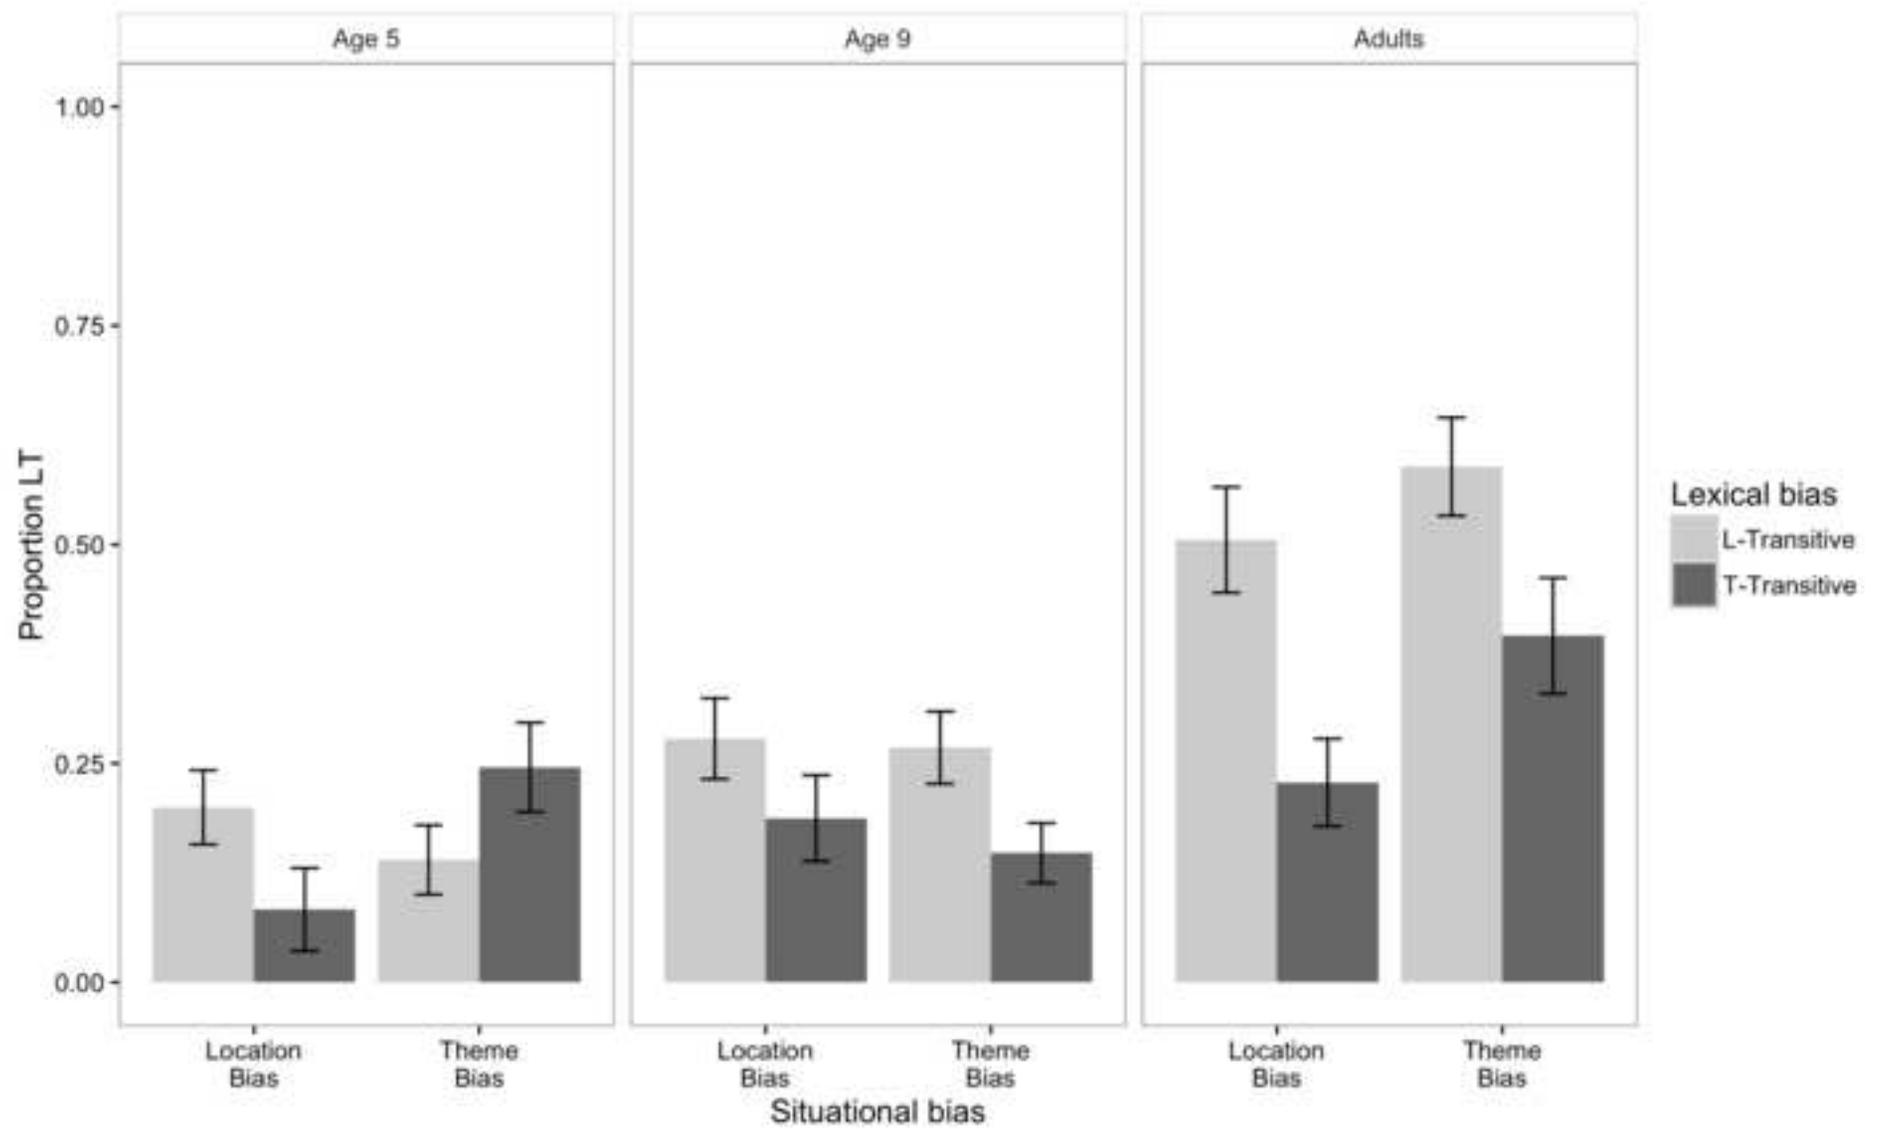


Figure8
[Click here to download high resolution image](#)

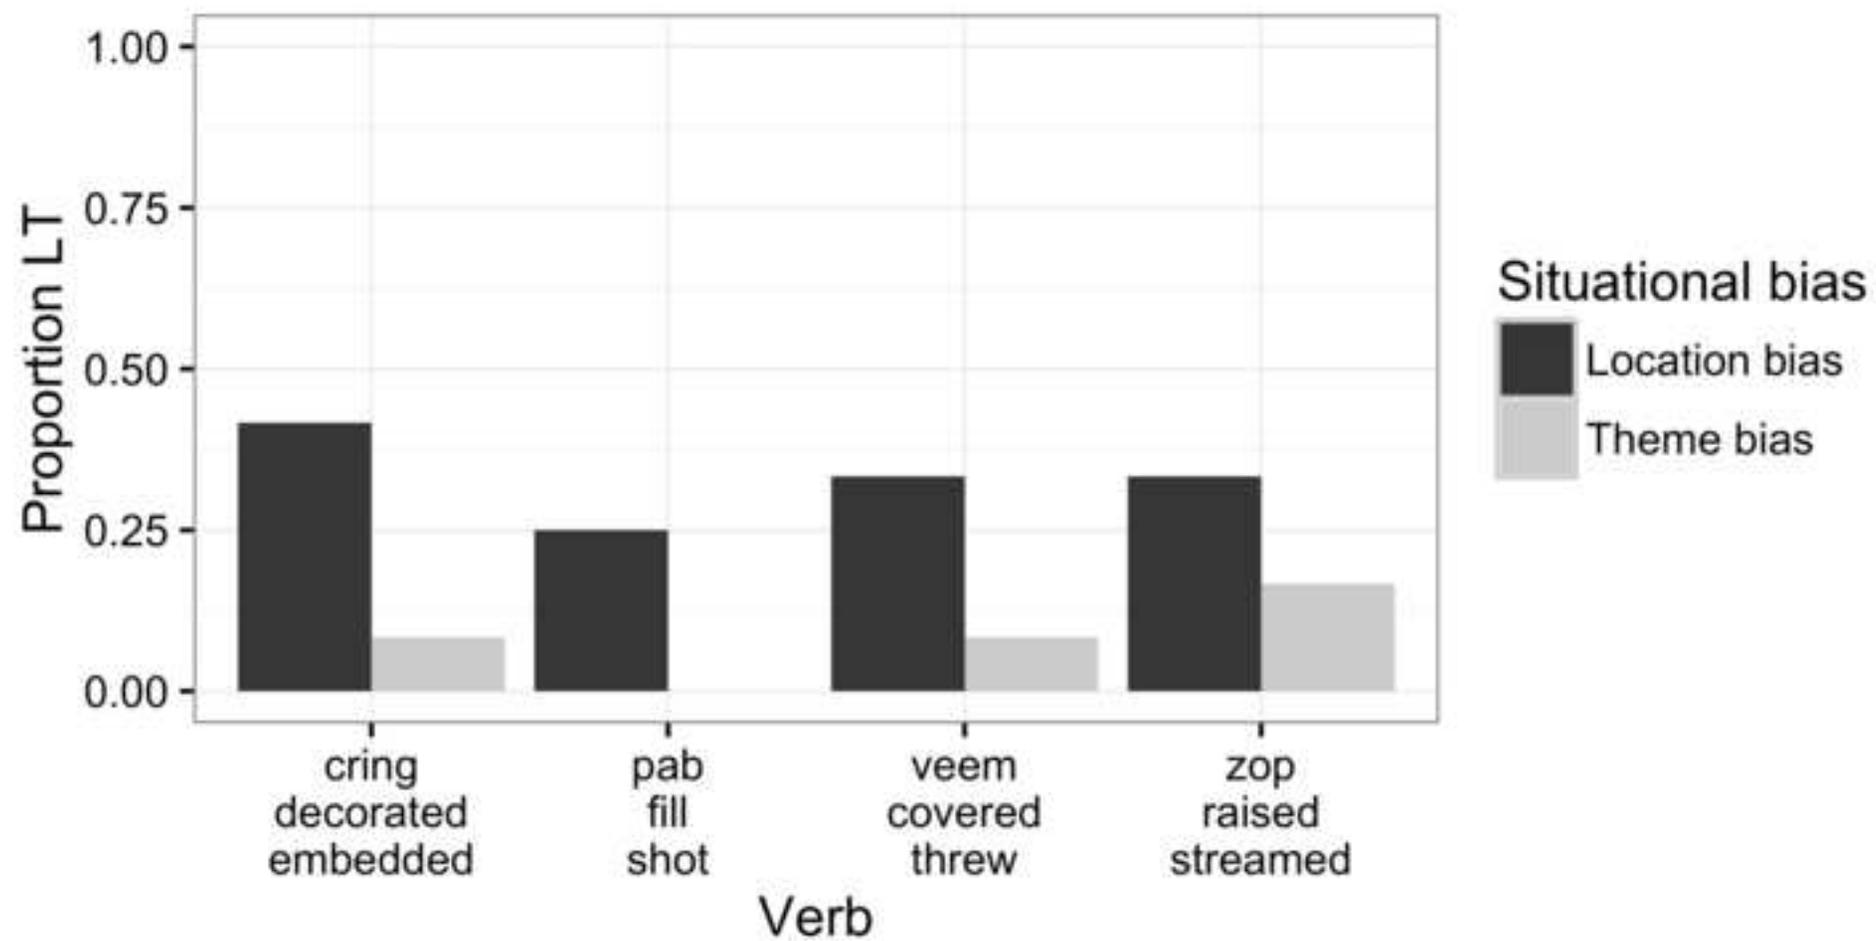


Figure9

[Click here to download high resolution image](#)

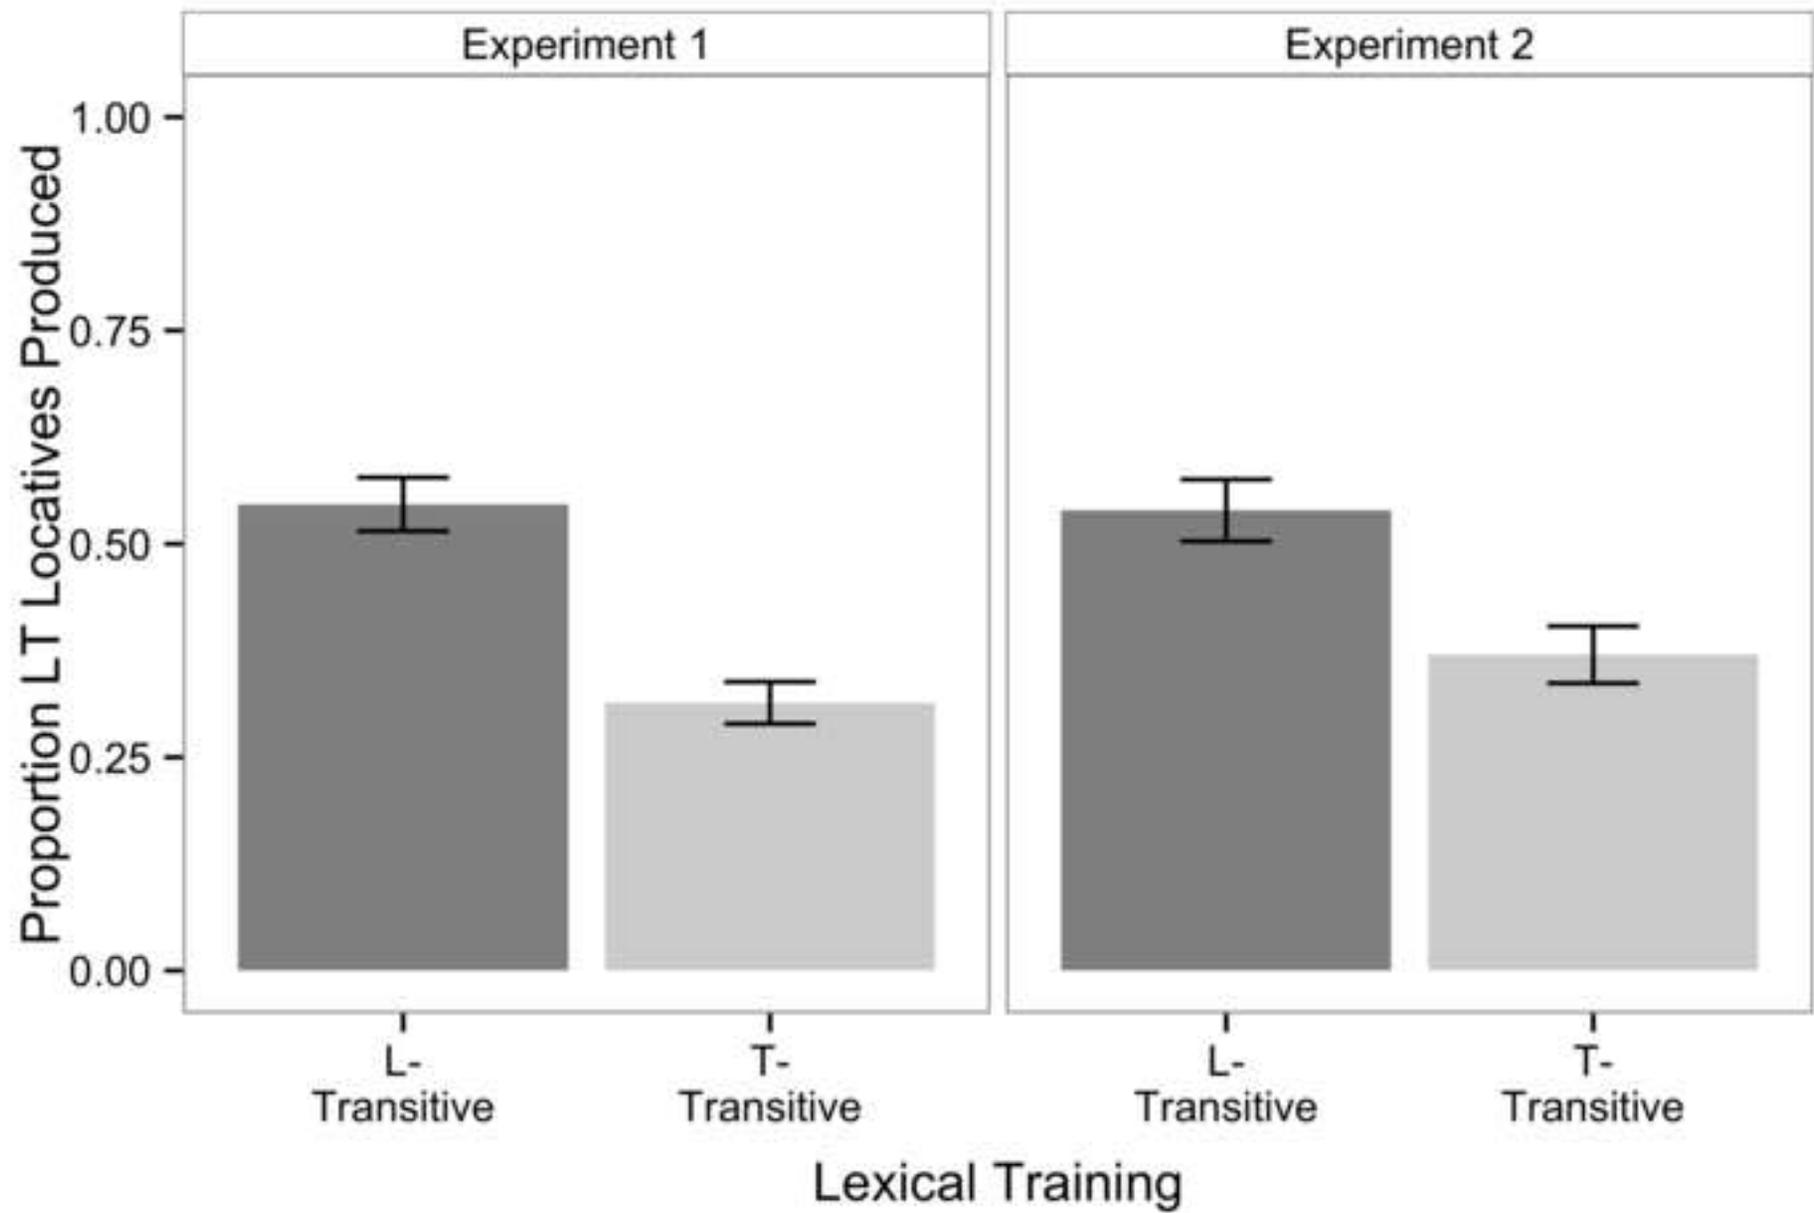


Figure10
[Click here to download high resolution image](#)

