Routledge
Taylor & Francis Group

Check for updates

REVIEW ARTICLE

What have labels ever done for us? The linguistic shortcut in conceptual processing

Louise Connell

Department of Psychology, Lancaster University, Lancaster, UK

ABSTRACT
How does language affect cognition? Is it important that most of our concepts come with linguistic labels, such as *car* or *number*? The statistical distributions of how such labels co-occur in language offers a rich medium of associative information that can support conceptual processing in a number of ways. In this article, I argue that the role of language in conceptual processing goes far beyond mere support, and that language is as fundamental and intrinsic a part of conceptual processing as sensorimotor-affective simulations. In particular, because linguistic association tends to be computationally cheaper than simulation (i.e. faster, less effortful, but still information-rich), it enables an heuristic mechanism that can provide adequate conceptual representation without the need to develop a detailed simulation. I review the evidence for this key mechanism – the linguistic shortcut – and propose that it allows labels to sometimes carry the burden of conceptual processing by acting in place of simulated referent meanings, according to context, available resources, and processing goals.

All right, but apart from the sanitation, the medicine, education, wine, public order, irrigation, roads, a fresh water system, and public health, what have the Romans ever done for us? (*Life of Brian*: Goldstone & Jones, 1979)

A concept in long-term memory is an aggregate of experience that receives frequent attention and can be re-activated relatively easily in offline processing as an instantiated representation of the concept[1] (e.g. Connell & Lynott, 2014b). We mentally represent things that are present in the environment during online processing of real-world perception and action, and can represent them again in their absence during offline remembering, planning, and day-dreaming (Wilson, 2002). As such, it seems reasonable that attaching a linguistic label to a particular aspect of experience could make it easier to perform some conceptual tasks, such as acquiring concepts like object kinds (Xu, 2002) and numbers (Carey, 2004), shaping the boundaries between colour categories (Winawer et al., 2007), or influencing how easily a visual feature or object is detected (Lupyan & Ward, 2013). Attaching a label to a bundle of experience may allow us to attend to it more easily in online processing, re-activate it more easily in offline processing, and hence help it cohere into a concept by assisting with abstraction (i.e. moving from a specific instance of sensorimotor experience to a generalised, aggregate form). However, the

possibilities regarding language's role in cognition are more far-reaching than assisting concept learning or online perception. Rather than being peripheral to the "real" concept at hand and playing a supporting role in cognition, language could instead be an integral part of the human conceptual system (Connell & Lynott, 2014b).

That is, the labels we attach to frequently-attended bundles of experience could comprise an indispensable part of concepts, play an active role in virtually all aspects of conceptual processing, and ultimately enable a more efficient form of cognition than would be possible without language. In the present article, I concentrate on an especially valuable form of information from linguistic labels, that of statistical, distributional associations, and review the evidence for its role in conceptual processing. In particular, I hypothesise how one key mechanism – the linguistic shortcut – allows labels to carry the burden of conceptual processing under a range of circumstances by effectively acting in place of deeper, more detailed representations of referent meaning. When it comes to cognition, labels do rather a lot for us.

Concepts as simulated and linguistic information

Many researchers from across the cognitive sciences have come to an interdisciplinary consensus that the

human conceptual system comprises two different types of information: simulated and linguistic (Barsalou, Santos, Simmons, & Wilson, 2008; Connell & Lynott, 2014b; Louwerse & Jeuniaux, 2008; Vigliocco, Meteyard, Andrews, & Kousta, 2009). Simulated information (also known as grounded or embodied representations) emerges from our interactive experience with the world around us, and represents conceptual knowledge as a simulation (i.e. partial replay) in the same neural systems that are active in processing real–world perception and action (Allport, 1985; Barsalou, 1999; Connell & Lynott, 2014b; Glenberg & Gallese, 2012; Yee, Chrysikou, Hoffman, & Thompson-Schill, 2013). Although often described as "sensorimotor" for reasons of history and convenience, such representations are not confined to simple sensory and motor information and are also assumed to incorporate affective information (e.g. valence, arousal, and emotional states: Niedenthal, 2007) and information about the broader situational context (e.g. environmental, social, and intentional factors: Barsalou, 2003; Niedenthal, Barsalou, Winkielman, Krauth-Gruber, & Ric, 2005). A simulated representation of a *car*, for example, may include its silver colour (visual), the whirr of its engine (auditory), its soft upholstery (tactile), turning the steering wheel (hand/arm action), enjoyment of the driving experience (affective), and streets and other cars outside (situational).

Evidence for simulations includes neuroimaging and patient studies that have found shared activation between areas involved in modality–specific perceptual or action experience and their equivalents in conceptual processing. For instance, reading a sound-related word like "thunder" activates the auditory association cortex, and aphasic patients with atrophy of the auditory association cortex have impaired processing of such words (Bonner & Grossman, 2012; see also Boulenger et al., 2008; Goldberg, Perfetti, & Schneider, 2006; Hauk, Johnsrude, & Pulvermüller, 2004). Across a range of paradigms, behavioural experimentation has also revealed complex interactions between the representations required for processing the environment and those simulated for conceptual processing (e.g. Connell & Lynott, 2010, 2012, 2014a; Dils & Boroditsky, 2010; Zwaan & Taylor, 2006). For example, people are faster to make size judgements about manipulable objects when they are holding a ball between the hands, due to the importance of haptic perception in experiencing the size of *apples* and *grapes* (Connell, Lynott, & Dreyer, 2012). The same effect does not occur for non-manipulable object like *elephant* or *deer*, where size experience and judgements tend to rely on visual information. In short, the conceptual system has effectively co-opted the perceptual, motor, and affective systems for the purposes of

representation (Connell & Lynott, 2014b), and simulated information is central to much of human conceptual processing.

The importance of simulated information in mental representation, however, does not mean that it alone subserves conceptual processing. Linguistic information emerges from our experience with language, and represents conceptual knowledge as statistical patterns of how words are distributed in relation to one another (Barsalou et al., 2008; Connell & Lynott, 2014b; Louwerse, 2011; Lynott & Connell, 2010). Natural languages are full of statistical regularities: words and phrases tend to occur repeatedly in similar contexts, and sensitivity to such regularities provides a powerful generalised learning mechanism from early infancy (Aslin & Newport, 2012; Saffran, 2003; Smith & Yu, 2008). Repeated exposure to language – and a typical adult is exposed to approximately 100,000 words per day (Bohn & Short, 2009) – allows a complex web of word–to–word (and word–to–phrase, phrase–to–phrase, etc.) associations to develop. A linguistic representation of *car*, for instance, may include associated words and phrases that frequently appear in proximity to the word "car", such as "sports", "insurance", "park", "used", and "driver".

Information from language alone is powerful enough to capture many aspects human experience. Empirical evidence for its utility in conceptual processing has come from a range of tasks, including property verification (Louwerse & Connell, 2011), property generation (Santos, Chaigneau, Simmons, & Barsalou, 2011; Simmons, Hamann, Harenski, Hu, & Barsalou, 2008), conceptual combination (Connell & Lynott, 2013), semantic relatedness and spatial iconicity (Louwerse & Hutchinson, 2012; Louwerse & Jeuniaux, 2010), metaphor comprehension (Hutchinson & Louwerse, 2013; Liu, Connell, & Lynott, 2018), and SNARC effects (Hutchinson & Louwerse, 2014). For example, the order of word distribution in English is highly directional (e.g. "root" tends to be mentioned before "branch" more often than *vice versa*). Louwerse and Jeuniaux (2010) showed that people are influenced by this directional linguistic information when making judgements about the spatial location of objects. When asked whether a "branch" typically occurs above a "root", both speed and accuracy of responses were influenced by linguistic information about word order, even when this word order was inconsistent with spatial configurations in the real world (i.e. "root" … "branch" is the typical word order, even though *branch* … *root* is the typical spatial configuration). Critically, the same effect of word order appeared when pictures were used rather than words (e.g. presenting an image of a *branch* above a *root*). That is, the role of linguistic distributional information in conceptual

processing is not confined to language stimuli, but also influences ostensibly non-linguistic tasks with image stimuli, and thus has broad importance across human conceptual processing.

In theoretical terms, there is some disagreement as to the extent to which concepts rely upon linguistic versus simulated information. Barsalou et al.'s (2008) Language As Situated Simulation (LASS) theory favours simulation as the main driver of conceptual processing; linguistic information – however useful it may be – lends itself mainly to a subset of circumstances where superficial strategies can adequately support task performance. By contrast, Louwerse's (2011; see also Louwerse & Jeuniaux, 2008) Symbol Interdependency Hypothesis (SIH) accords a dominant role to linguistic rather than simulated information, and proposes that, due to the way it encodes perceptual relations in the real world, it is capable of doing much of the heavy lifting in a range of cognitive tasks. An alternative possibility lies in between these two accounts, namely that neither linguistic nor simulated information is fundamentally more important than the other across conceptual processing as a whole, but rather that the relative importance of each at a particular point in time in a particular individual depends on a variety of factors. As detailed elsewhere (Connell & Lynott, 2014b, TopiCS), a conceptual representation is in itself a process, whereby distributed neural activation patterns across sensorimotor, affective, language, and other areas coordinate to subserve the requirements of the moment, and thereby include both simulated and linguistic information. So sensitive is this process to factors including concurrent sensorimotor and attentional processing, the availability of appropriate labels, and the accumulation and retrieval of sensorimotor and linguistic experience over time, that even if one attempts to hold the task requirements constant, an individual's representation of a particular concept (e.g. car) is never precisely the same from one occasion to the next. Neither simulated nor linguistic information necessarily dominate such conceptual representations because both are needed to provide sufficient flexibility and robustness to a cognitive system that must operate continuously in a noisy environment with limited resources.

Is linguistic information grounded?

Although its basic form is the label, linguistic information is not amodal in that it does not comprise ungrounded symbols. When considering the mechanics of grounding, however, it is important to distinguish between grounding of concepts in long-term memory and grounding of a particular instantiated representation. A concept in long-term memory is an aggregate of sensory, motor, affective, situational, and linguistic (i.e. label) experience, that develops by repeatedly attending to aspects of experience that have overlapping patterns of neural activation. Linguistic information in long-term memory is therefore grounded because labels are attached to the sensorimotor-affective-situational experience of their referent. However, the instantiated representation of a concept is a dynamic and transient entity, constructed on the fly according to the constraints of task demands, context, available resources, and processing goals, and includes perceptual, motor, affective, situational, and linguistic (i.e. label) information to greater or lesser degrees (Connell & Lynott, 2014b). Linguistic information in a given instance of conceptual representation is therefore not necessarily grounded *at that particular point in time* because the specific information that makes up a dynamic mental representation will vary and may include little or no simulated information alongside the label. That is, "the concept to which a word refers is ultimately grounded in the simulation system; however, a word does not need to be fully grounded every time it is processed" (Louwerse & Connell, 2011, p. 393). This notion of limited grounding is central to the SIH (Louwerse, 2011; Louwerse & Jeuniaux, 2008), which argues that, because labels link to each other as well as to grounded referents, it is possible for language comprehenders to ground only some labels and then bootstrap the meaning of others through distributional patterns between the labels themselves.

For instance, imagine someone develops their concept of *car* in long-term memory via a wide range of perceptual, motor, and affective experiences of cars across multiple situational contexts, and via a wide range of related words in statistical distributions around the label "car" across multiple contexts. The long-term concept of *car* therefore has a rich variety of information potentially available, and a particular instantiation of *car* will draw upon a subset of this information to create an appropriate representation for the demands and constraints of the occasion. Sometimes, a particular instantiation of *car* may include little perceptual, motor, or affective information and may instead rely heavily on linguistic information: for instance, when judging rapidly that "car" and "park" are semantically related. Thus, although the concept *car* is indeed fully grounded in terms of the information available in long-term memory, all this grounded information is not necessarily represented during a given instance of conceptual processing.

In summary, concepts in long-term memory comprise both linguistic and simulated information, and so the labels that make up linguistic information are grounded

in the sensorimotor-affective-situational experience that makes up simulated information. Because the content of instantiated representations varies to fulfil momentary demands, however, labels are not necessarily grounded on every occasion of conceptual processing.

Does linguistic information simply mirror simulated information?

Linguistic distributional statistics and simulated distributional statistics contain similar patterns, but do not directly reflect one another. In contrast to linguistic information, which comprises statistical regularities between word forms, simulated information encodes statistical regularities at the level of meaning due to the inclusion of situational context in simulated representations. A *car*, for instance, typically has *wheels* and a *driver*, operates on the *road* or *street*, and sometimes needs a *service* or *repair*. Objects, events, and other situational entities tend to occur together in the real world in ways that, through cumulative interactive experience, can give rise to statistical patterns of how referent concepts are distributed in relation to one another.

It might seem reasonable to expect language to mirror this real-world distribution of referent concepts. Indeed, both LASS (Barsalou et al., 2008) and SIH (Louwerse, 2011) implicitly assume that linguistic distributional statistics largely, if not completely, reflect the same information as sensorimotor experience of the real world. However, we do not talk about the world merely to describe it, but to question, analyse, interpret, abstract, and predict it. As such, I propose that the statistical distributions of how words co-occur with one another offers a rich medium of associative information that goes beyond the statistical distributions of how referent concepts co-occur in real-world experience. Linguistic information can therefore, in principle, capture qualitatively different aspects of conceptual representation to that which can be captured by simulated information.

Take, for instance, the concept of *democracy*. One learns about *democracy* by reading, hearing, and using its label alongside other words such as "government", "human rights", "election", "freedom", "vote", and so on. One also learns about democracy by going to a polling station and marking a ballot in order to cast a vote. While these linguistic and sensorimotor experiences of *democracy* are related, they differ in their broader situational contexts. Sensorimotor experience (real or vicarious) is limited to physical interactions with concrete entities and there is no direct concrete referent of *democracy*, only of related concepts such as *polling station*, *ballot paper*, and so on. Of course, the experience of

reading, hearing, and speaking about *democracy* also has a sensorimotor basis, but it is concerned with the visual, auditory, and motor experience of word forms, and not with sensorimotor experience of the words' referents. Hence, the statistical patterns of how entities and events are distributed in sensorimotor experience of *democracy* will not easily capture its relationship with concepts such as *freedom* or *human rights*. By contrast, linguistic experience can include both descriptions of sensorimotor experience of a label's referent *and* metaphysical discussions of the concept itself. The experience of the label "democracy" can include contexts where its history is analysed, its validity questioned, its effects on society interpreted, and other high-level topics of discussion that have no direct sensorimotor correlate. That said, linguistic experience is limited to situations that people choose to speak and write about and so cannot perfectly describe all sensorimotor experience. The statistical patterns of how words are distributed in linguistic experience of *democracy* therefore cannot easily capture certain aspects of a real-world experience, such as those that are relatively unimportant (e.g. the precise colour of the pencil in a polling booth), or highly personal or taboo (e.g. an individual's unspoken aversion to a candidate's name).

To summarise, linguistic distributional statistics are not a deterministic function of sensorimotor distributional statistics, and – although they share mutual information – each should be capable of contributing unique information to conceptual processing. While sensorimotor experience can capture a richness and precision of detail that remains unspoken and therefore eludes linguistic information, the statistical regularities in language have, in principle, an ability to inform us about conceptual relations in a way that the statistical regularities in sensorimotor-affective experience of the real world cannot. There is some empirical support for this position from computational modelling research. For example, Louwerse and Connell (2011) examined whether linguistic distributional information could predict the perceptual modality of various sensory adjectives (e.g. *rustling, speckled, delicious*: see also Lynott & Connell, 2009, 2013). Using co-occurrence frequencies from a large corpus to model word distributions, they found that words could be successfully classified into auditory, visual-haptic, and olfactory-gustatory clusters with a high degree of accuracy. However, it was not possible to separate touch-related words like *sharp* from sight-related words like *speckled*, nor taste-related words like *delicious* from smell-related words like *smoky*, on the basis of language alone. Such evidence suggests that the structure of linguistic distributional information sometimes approximates that of simulated information but is less detailed and

precise concerning sensorimotor details, meaning that simulated information is uniquely responsible for sensorimotor precision (see also Barsalou et al., 2008; Connell & Lynott, 2014b; Riordan & Jones, 2011). On the other hand, Andrews, Vigliocco, and Vinson (2009) demonstrated that linguistic information also has the capacity to capture aspects of experience that simulated sensorimotor information may not. They compared the structure of distributional statistics derived from corpus analysis (i.e. linguistic information) and feature-listing norms (i.e. simulated information). They found that linguistic distributions tended to emphasise more encyclopedic information (e.g. taxonomies and abstractions: *eat → food*, *diet*, *eaten*, *cereals*) whereas feature distributions tended to emphasise more sensorimotor information (e.g. physical acts and interactions: *eat → vomit*, *taste*, *teeth*, *drink*). Moreover, the combination of linguistic and simulated information was better able to predict semantic priming data than either type of information alone. Linguistic distributional information can therefore do more than simply – and imperfectly – reflect sensorimotor experience, and appears to be capable of capturing a qualitatively different form of information that makes a unique contribution to conceptual processing.

While the above evidence supports the general principle, a number of key questions remain unanswered regarding the extent to which the statistical regularities in language go beyond a mere reflection of the statistical regularities in sensorimotor-affective experience of the real world. For example, what proportion of information in linguistic distributional statistics is unique? And is this unique information systematically restricted to certain types of concepts or randomly distributed throughout the conceptual system? Several theories of conceptual representation now assume that abstract concepts are acquired via language, and are represented via linguistic information, to a greater extent than are concrete concepts (Borghi & Binkofski, 2014; Crutch & Warrington, 2005; Vigliocco et al., 2009). Linguistic information might be relatively more important to an abstract concept like *democracy*, for example, compared to a concrete concept like *chair*. Moreover, it appears that many, if not most, of our concepts might be abstract rather than concrete; in Brysbaert, Warriner, and Kuperman's (2014) concreteness norms for words that are known by at least 85% of English speakers, some 53% of 40,000 words fell in the abstract end of the scale. It is therefore possible that any unique information from linguistic distributional statistics is likely to be present to a greater extent in abstract than in concrete concepts, and – given the large number of abstract concepts in our conceptual system – that the proportion of unique information is far from negligible.

What advantage has linguistic information?

One critical difference between processing simulated versus linguistic information is that linguistic association is typically faster than simulation (Barsalou et al., 2008; Connell & Lynott, 2014b; Lynott & Connell, 2010). When a word is presented, activation of linguistic associates reaches peak activation before relatively slower simulated information is fully available (see LASS theory: Barsalou et al., 2008). This speed advantage is a general tendency rather than an absolute difference, and may vary with the demands and constraints of the wider situation. Nonetheless, supporting evidence has been found in a range of paradigms, including behavioural (Connell & Lynott, 2013; Louwerse & Connell, 2011; Santos et al., 2011), electrophysiological (Louwerse & Hutchinson, 2012), and neuroimaging (Simmons et al., 2008) studies. For instance, when asked to list properties for a given concept, Santos et al. (2011) found that people tended to begin by listing linguistic associates of the word (e.g. *bee → hive*, *honey*, *sting*) and later moved onto non-associates that were consistent with simulating the concept in a broad situational context (*bee → wings*, *summer*, *flowers*). Moreover, neural activation while listing the early properties overlapped with individually-established activations for word association (Broca's area in the left inferior frontal gyrus), and activation while listing the late properties overlapped with activations for deliberately imagining a situation (right posterior superior temporal sulcus: Simmons et al., 2008).

However, the speed advantage of linguistic association does not mean that sensorimotor-affective simulation is slow. Simulated information is activated extremely rapidly. In word reading, for instance, action words relating to the leg versus face produce topographic differences in neurophysiological activation approximately 200 ms after word onset (Hauk & Pulvermüller, 2004), suggesting very rapid access to effector-specific motor simulations. Nonetheless, sensitivity to linguistic distributional information also occurs very early. Contexts that frequently versus infrequently co-occur with a target word elicit greater negativity in distributed/frontal regions approximately 120 ms after target word onset, even when both contexts are semantically related to the target with similar cloze probability (Molinaro, Barraza, & Carreiras, 2013), which suggests very rapid access to common linguistic associates.

It is essential, however, not to conflate early neural activation with conceptual activation at task completion. The crux of the linguistic speed advantage has more to do with the relative activation of conceptual information that can inform a response (i.e. at what point in the time-course of a given task does the word "car" activate

critically useful information?), than with literal neurophysiological activation (i.e. at what point can the word "car" produce a differential activation trace in EEG/MEG?). Indeed, because the systems of linguistic and simulated information interact and reinforce one another in cycles of activation during conceptual processing (i.e. linguistic information can activate simulated information, which in turn can activate further linguistic information, and so on: Barsalou et al., 2008; Connell & Lynott, 2014b), early neural activation in one system may lead to a response that is primarily driven by the other system. That is, there is a disconnect between the question of whether linguistic or simulation areas first show measurable neurophysiological activation, and the question of whether the conceptual activation that drives a particular response is primarily linguistic or simulated in nature. For example, take a semantic priming task where the prime word is "car" and the target word is "driver". The word "driver" can be conceptually activated in two ways: via linguistic association from the prime "car", or via labelling the sensorimotor-affective simulation of *car with a driver behind the steering wheel*. When "car" primes "driver", which form of conceptual activation is responsible? At this point, it does not particularly matter whether the sensorimotor cortex or linguistic association areas were activated first on presentation of the word "car". Rather, what matters is the relative activation of linguistic and simulated information at the point that a response is required. In our example, linguistic association from "car" → "driver" will typically be faster than simulating a situation and then labelling it, as in "car" → *car with a driver behind the steering wheel* → "driver", and so linguistic information will primarily drive the response to the target word "driver". Other responses may be primarily driven by simulated information if the target word is not a close linguistic associate of the prime, but the critical point remains. When it comes to conceptual activation, linguistic association is typically faster than simulation.

The linguistic shortcut hypothesis

To recap, concept labels do not have to be grounded every time they are processed, the distributional patterns of labels have the capacity to capture unique conceptual information as well as approximate sensorimotor experience, and label-to-label association operates faster than simulation. Together, these characteristics mean that linguistic associates can be viewed as computationally cheaper than simulation – faster, less effortful, but still information-rich – which enables an heuristic mechanism in conceptual processing termed the *linguistic shortcut* (Connell & Lynott, 2013, 2014b; Lynott & Connell,

2010). If computationally cheaper information from linguistic associates can usefully inform a response in a particular task before relatively more expensive simulated information can do so, then linguistic information effectively has the potential to act as a shortcut during conceptual processing. The main impact of this linguistic shortcut is that associative information between concept labels can sometimes provide adequate conceptual representation without the need to develop a detailed simulation. In other words, a label can, at times, act in place of sensorimotor simulation of its referent meaning.

In the above example of processing the word "car", activating the label "driver" via linguistic association is not necessarily any use to the particular task at hand; whether processing can stop here, or whether further processing must (and can) be pursued in the simulation and/or linguistic system, varies according to circumstance. Conceptual processing does not happen in a vacuum but usually occurs with the goal of making a response, whether it's a specific response in an experimental paradigm (e.g. "Peach can be fuzzy?": press "yes" or "no" button), a spoken responses in real-world dialogue (e.g. "What's in a Bellini?" → "Prosecco and peach juice", or "I don't know"), an action response to environmental stimuli (e.g. "Are there any peaches left?" → find a *peach* in a fruit bowl full of apples, apricots, and nectarines), or simply a mental response in an ongoing train of thought (e.g. What do you call that fruit that's like a peach but not fuzzy? → "nectarine"). When and where the linguistic shortcut is likely to be used to inform a response in a conceptual task depends on many factors, including depth of processing demands, the resources available for processing, and motivation of the individual. The rest of this section reviews the evidence for several such circumstances.

Depth of processing demands

The linguistic shortcut becomes particularly useful when a task allows relatively shallow or superficial conceptual processing. Shallow processing in this sense means that each word (in a sentence, paragraph, or other task) does not contribute its full potential meaning and/or is integrated incompletely into the broader context, resulting in an underspecified representation (Sanford & Sturt, 2002). That is, not every task requires a wholly accurate and fully developed conceptual representation, and sometimes a fuzzy, "good enough" approximation will suffice (Ferreira, Bailey, & Ferraro, 2002). Deep processing, on the other hand, means each word contributes a full and detailed meaning and is integrated completely into the broader context,

resulting in a fully-specified representation. The difference between shallow and deep processing is relative and perhaps best illustrated with an example. When asked to answer the question "When an aircraft crashes, where should the survivors be buried?", only 20% of participants noticed the anomaly that *survivors* are living people whom one does not bury (Barton & Sanford, 1993). That is, most people have processed the word "survivors" only shallowly because it fits the situation described in the rest of the sentence and does not jump out as meriting closer attention. When the question asked about a bicycle accident instead of an air crash, detection rates rose to 80% because the word "survivors" does not fit and led most people to process its meaning more deeply.

The representation of word meaning for a given individual in a given context depends on what particular information makes up the dynamic mental representation at hand, and so long as an ongoing representation can appropriately fulfil current task goals, it is fit for purpose (Connell & Lynott, 2014b). Both the LASS theory (Barsalou et al., 2008) and SIH (Louwerse & Jeuniaux, 2008; see also Louwerse & Connell, 2011) hold that linguistic information is well suited to shallow conceptual processing where a "good enough" approximation will suffice, whereas simulated information is required for deeper conceptual processing where precise detail and/or integration into the broader situation is necessary. One would therefore expect that, unless a task requires deep and precise conceptual processing, a concept might be represented and processed via its label and associated linguistic information rather than via a detailed simulation of its referent. In support, Solomon and Barsalou (2004) found that people tended to rely on linguistic association in a property verification task (e.g. true or false: *salmon→scales*) when the false filler items were unassociated (e.g. *bicycle→chin*), because the conflation of true/false with associated/unassociated allowed a shallow association strategy to suffice. Only when difficult associated false fillers were included (e.g. *banana→monkey*) did people begin to process the concepts more deeply, and linguistic association could no longer predict responses.

Connell and Lynott (2013) took the depth of processing proposal one step further by examining the same stimuli in two related tasks: shallow sensibility judgement of novel noun-noun phrases (e.g. yes or no: does *elephant complaint* make sense?) and deep interpretation generation (e.g. yes or no: can you think of a meaning for *elephant complaint*?). Critically, although the stimuli were all novel, the constituent nouns varied in their distributional co-occurrence frequencies. They found that linguistic distributional information predicted

the speed of accepting a phrase as sensible but not accepting a phrase as interpretable. That is, sensibility judgement could get away with using the linguistic shortcut as an heuristic: the more often the constituent words appeared together in language, the more quickly people accepted the phrase as making sense. In contrast, interpretation generation – that is, actually coming up with a meaning – needed deeper processing that linguistic information could not offer, and presumably relied on simulated information to generate a meaning (see also Liu et al., 2018).

Word meaning can potentially be represented anywhere on the continuum between extremely shallow linguistic association and extremely deep sensorimotor-affective-situational simulation, and any point on this continuum is as valid a form of meaning as any other if it allows someone to read a novel, communicate in dialogue, press buttons in an experiment, or whatever the response goal might be. By exploiting the distributional information in label-to-label associations, the linguistic shortcut offers a quick heuristic that may be good enough to inform a response without requiring deeper processing. Where possible, when a shallow response will suffice, a label may act in place of a detailed simulation of its referent.

Available processing resources

The linguistic shortcut is also useful when a task begins to strain available cognitive resources. While simulating sensorimotor experience of a familiar concept such as *car* or *banana* might be relatively straightforward, it comprises only a small part of conceptual processing. People must also be able to represent new concepts, process familiar concepts in new ways, and do so within the constraint that such conceptual processing might not always be successful. Humans do not have infinite executive and memory resources to devote to conceptual processing, and so it would greatly enhance cognitive efficiency if there were some form of triage mechanism available that could flag up at an early stage when a particular conceptual process is likely to prove excessively difficult and is not worth further effort.

Connell and Lynott (2013; Lynott & Connell, 2010) propose that the linguistic shortcut can offer such a cognitive triage mechanism. Since concepts that are often discussed in the same context are likely to be easily integrated in a shared situation (Lynott & Connell, 2010), linguistic distributional information from concept labels offers a guide to the likelihood of successful simulation. Moreover, since linguistic information is available quickly and computationally cheaply, it can provide a "quick and dirty" heuristic to determine whether it is

worth expending precious cognitive effort on a particular conceptual processing task, or whether such processing should be abandoned pending further clarification or information. That is, conceptual processing can be halted at the point of representing concepts via their labels and associated linguistic information rather than continuing to develop a detailed simulation that may fail to cohere or otherwise incur unnecessary costs.

The idea of the linguistic shortcut acting as a cognitive triage mechanism has empirical support. As previously discussed, Connell and Lynott (2013) asked people to either judge the sensibility or provide an interpretation of novel noun-noun compounds (e.g. *elephant complaint*): a conceptual combination task that required creating a new representation. However, Connell and Lynott did not only examine "yes" responses that represented successful conceptual processing (i.e. responses where the compounds were judged as sensible or an interpretation was provided), but also "no" responses that would traditionally be discarded as representing unsuccessful processing of the task (i.e. responses where compounds were deemed to make no sense or have no meaningful interpretation). They found that people were more likely to reject a novel conceptual combination as nonsensical or uninterpretable, and do so more rapidly, when the constituent words rarely co-occurred across language. That is, linguistic distributional information predicted both the likelihood and speed of abandoning conceptual processing; exactly as expected from a cognitive triage mechanism.

Liu et al. (2018) explored the mechanism in more detail by introducing a competing measure of ease of simulation (i.e. a normed variable that reflects how easily a phrase could result in a coherent simulation of meaning) and comparing its effect to that of distributional co-occurrence frequency in metaphor processing. Following Connell and Lynott (2013), Liu et al. found that people were more likely to reject a sentence as nonsensical when the constituent words seldom co-occurred across language (e.g. *illness can be bright*) compared to when they co-occurred more often (e.g. *supply can be bright*), and additionally showed that this effect occurred independently of ease of simulation. The speed of rejection, however, was more complex. People quickly rejected metaphors that were typically considered to be difficult to simulate, but this effect was moderated by co-occurrence frequency. When the words in a metaphor rarely co-occurred, ease of simulation had a reduced effect, suggesting that people rejected metaphors as nonsensical based on their low co-occurrence frequency without further processing of simulated information. In other words, the linguistic shortcut was used as a cognitive triage mechanism to identify processing

that was unlikely to result in a coherent simulation and should therefore be halted rather than potentially waste resources on unnecessary processing.

Because there is an inherent tension between the needs of conceptual processing to represent new concepts or word meanings and the limited cognitive resources that are available to do so, a system of cognitive triage can help manage the strain on resources by allowing processing to be abandoned, postponed, or prioritised according to the requirements of the moment. The linguistic shortcut offers one such cognitive triage mechanism by using the distributional information in label-to-label associations as a quick heuristic to flag up when a particular occasion of conceptual processing is likely to lead to unnecessary costs. That is, a label may act in place of of a detailed simulation of its referent when it offers the possibility to optimise or satisfice conceptual processing in a limited-resource cognitive system.

Conclusions

So apart from the approximation of sensorimotor information, the contribution of unique information, optional grounding, enhanced speed, a computationally cheap linguistic shortcut, a means of "good enough" processing, and a cognitive triage mechanism, what have labels ever done for us? As future research will most likely demonstrate, quite a bit. We are only at the beginning of establishing exactly how, where, and when conceptual processing relies on labels and their associations. A wide variety of evidence underscores the important role that linguistic information plays in cognition and conceptual processing, but there remain several outstanding questions about the circumstances that affect the relative roles of linguistic versus simulated information. For instance, the content of mental representations dynamically changes with available resources, current goals, and individual motivations. Does reliance on computationally cheap linguistic information therefore increase when available resources are restricted, such as when participants must respond under time pressure, carry out a concurrent cognitive task, or retain a load in memory? Or when participant motivation is low, such as when people are unconcerned about performing the task accurately due to fatigue, boredom, or lack of reward? More research is needed to address these issues in order to pin down a proper understanding of how linguistic labels enable fault tolerance in conceptual processing, and allow cognition to respond dynamically and flexibly to varying demands and constraints.

The Linguistic Shortcut Hypothesis shares many commonalities with Barsalou et al. (2008) LASS theory and Louwerse's (2011) SIH, but there are three critical differences. Firstly, while LASS favours simulation as the main driver of conceptual processing and SIH favours linguistic information, the Linguistic Shortcut Hypothesis proposes a middle ground: neither linguistic nor simulated information is fundamentally more important than the other across conceptual processing as a whole, and their relative importance in a given response depends on a variety of factors. Secondly, both LASS and SIH assume that linguistic distributional statistics reflect (albeit imperfectly) the same information as sensorimotor experience of the real world, whereas I propose that linguistic information captures unique conceptual information and enables some conceptual processing that would not be possible with sensorimotor simulation alone. Both of these differences have been discussed earlier in the paper. There is, however, a third difference. LASS and SIH both describe how linguistic information can help conceptual processing reach a successful outcome: a response that fulfils the requirements of the task. That is, conceptual processing can proceed via slow-but-accurate simulation or fast-but-approximate linguistic information; sometimes linguistic information will suffice, but where it does not, conceptual processing can fall back on a detailed sensorimotor simulation. However, the Linguistic Shortcut Hypothesis goes further by proposing that linguistic information has a broader role that facilitates not only the successful outcome of conceptual processing (i.e. a response that fulfils task requirements) but also the unsuccessful outcome (i.e. a response that does not fulfil task requirements), and – critically – helps to determine which outcome should be pursued. While it may seem obvious that a successful outcome is desirable from the perspective of the task at hand, an unsuccessful outcome is also potentially desirable from the perspective of conserving costs in a cognitive system that has limited time and resources to devote to any particular task.

The cognitive triage mechanism of the Linguistic Shortcut Hypothesis uses linguistic information to guide which outcome should be pursued and how much effort should be spent in doing do. Sometimes linguistic information will suggest task requirements *cannot* be reasonably met, in which case conceptual processing can quickly and actively opt for an unsuccessful outcome and abandon simulation without expending further effort. Sometimes linguistic information is not particularly indicative of whether or not task requirements can be met, and so conceptual processing will continue to expend time and effort on simulation in pursuit of an uncertain outcome (i.e. a successful or unsuccessful outcome will depend on whether a detailed sensorimotor simulation can meet task requirements). Finally, sometimes linguistic information will suggest task requirements *should* be easily met, in which case conceptual processing will pursue a successful outcome as outlined by LASS and SIH: sometimes linguistic information will suffice, but where it does not, a detailed sensorimotor simulation can be used as a fallback option. To give the distinction an empirical focus, SIH and LASS can predict the latency of conceptual processing that ultimately "works" and fulfils the requirements of the given task (e.g. listing object features, Santos et al., 2011; verifying object properties: Louwerse & Connell, 2011; judging that a novel phrase makes sense: Connell & Lynott, 2013). The Linguistic Shortcut hypothesis can do that *and* can additionally predict the latency of conceptual processing that ultimately stalls or fails to meet the requirements of the task (e.g. deciding that one *cannot* think of a meaning for a novel phrase: Connell & Lynott, 2013; deciding a metaphor does *not* make sense: Liu et al., 2018). Failure to fulfil the requirements of a particular task is not necessarily an error, or the unfortunate result of conceptual processing gone wrong, but is sometimes a sensible and systematic option to conserve effort in a cognitive system of limited resources. Further research should explore the timecourse of unsuccessful conceptual processing as a means of establishing how linguistic versus simulated information contribute to cognitive efficiency.

Whorf (1956, p. 212) contended that "Language is not merely a reproducing instrument for voicing ideas but rather is itself the shaper of ideas, the programme and guide for the individual's mental activity". While the first part may be true – language is far more than a vehicle for meaning – the second part of Whorf's contention is too restrictive. Language is not a shaping *programme* for perception, action, and other aspects of cognition; such a programme would be unnecessarily narrow and imprison cognitive function in the labels and structure of language. Rather, linguistic labels are an essential springboard in the framework of human cognition that enables concepts to be represented and manipulated with more flexibility and efficiency than would otherwise be possible.

Note

1. This distinction between long-term concepts and instantiated representations follows the type-token distinction.

Disclosure statement

No potential conflict of interest was reported by the authors.

References

Allport, D. A. (1985). Distributed memory, modular subsystems and dysphasia. In S. K. Newman & R. Epstein (Eds.), *Current perspectives in dysphasia* (pp. 207–244). Edinburgh: Churchill Livingstone.

Andrews, M., Vigliocco, G., & Vinson, D. (2009). Integrating experiential and distributional data to learn semantic representations. *Psychological Review*, 116, 463–498.

Aslin, R. N., & Newport, E. L. (2012). Statistical learning from acquiring specific items to forming general rules. *Current Directions in Psychological Science*, 21, 170–176.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577–660.

Barsalou, L. W. (2003). Abstraction in perceptual symbol systems. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 358, 1177–1187.

Barsalou, L. W., Santos, A., Simmons, W. K., & Wilson, C. D. (2008). Language and simulation in conceptual processing. In M. De Vega, A. M. Glenberg, & A. C. Graesser (Eds.). *Symbols, embodiment, and meaning* (pp. 245–283). Oxford: Oxford University Press.

Barton, S. B., & Sanford, A. J. (1993). A case study of anomaly detection: Shallow semantic processing and cohesion establishment. *Memory & Cognition*, 21, 477–487.

Bohn, R. E., & Short, J. E. (2009). *How much information? 2009 report on American consumers*. San Diego, CA: UCSD Global Information Industry Center.

Bonner, M. F., & Grossman, M. (2012). Gray matter density of auditory association cortex relates to knowledge of sound concepts in primary progressive aphasia. *Journal of Neuroscience*, 32, 7986–7991.

Borghi, A. M., & Binkofski, F. (2014). *Words as social tools: An embodied view on abstract concepts*. New York, NY: Springer.

Boulenger, V., Mechtouff, L., Thobois, S., Broussolle, E., Jeannerod, M., & Nazir, T. A. (2008). Word processing in Parkinson's disease is impaired for action verbs but not for concrete nouns. *Neuropsychologia*, 46, 743–756.

Brysbaert, M., Warriner, A. B., & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*, 46, 904–911. http://dx.doi.org/10.3758/s13428-013-0403-5

Carey, S. (2004). Bootstrapping & the origin of concepts. *Daedalus*, 133, 59–68.

Connell, L., & Lynott, D. (2010). Look but don't touch: Tactile disadvantage in processing modality-specific words. *Cognition*, 115, 1–9.

Connell, L., & Lynott, D. (2012). When does perception facilitate or interfere with conceptual processing? The effect of attentional modulation. *Frontiers in Psychology*, 3(474), 1–3.

Connell, L., & Lynott, D. (2013). Flexible and fast: Linguistic shortcut affects both shallow and deep conceptual processing. *Psychonomic Bulletin & Review*, 20, 542–550.

Connell, L., & Lynott, D. (2014a). I see/hear what you mean: Semantic activation in visual word recognition depends on perceptual attention. *Journal of Experimental Psychology: General*, 143, 527–533.

Connell, L., & Lynott, D. (2014b). Principles of representation: Why you can't represent the same concept twice. *Topics in Cognitive Science*, 6, 390–406.

Connell, L., Lynott, D., & Dreyer, F. (2012). A functional role for modality-specific perceptual systems in conceptual representations. *PLOS ONE*, 7, e33321, 1–7.

Crutch, S. J., & Warrington, E. K. (2005). Abstract and concrete concepts have structurally different representational frameworks. *Brain*, 128, 615–627.

Dils, A. T., & Boroditsky, L. (2010). Visual motion aftereffect from understanding motion language. *Proceedings of the National Academy of Sciences*, 107, 16396–16400.

Ferreira, F., Bailey, K. G., & Ferraro, V. (2002). Good-enough representations in language comprehension. *Current Directions in Psychological Science*, 11, 11–15.

Glenberg, A. M., & Gallese, V. (2012). Action-based language: A theory of language acquisition, comprehension, and production. *Cortex*, 48, 905–922.

Goldberg, R. F., Perfetti, C. A., & Schneider, W. (2006). Perceptual knowledge retrieval activates sensory brain regions. *Journal of Neuroscience*, 26, 4917–4921.

Goldstone, J. (Producer), & Jones, T. (Director). (1979). *Monty Python's life of brian*. [Motion picture]. London, UK: HandMade Films and Python (Monty) Pictures Limited.

Hauk, O., Johnsrude, I., & Pulvermüller, F. (2004). Somatotopic representation of action words in human motor and premotor cortex. *Neuron*, 41, 301–307.

Hauk, O., & Pulvermüller, F. (2004). Neurophysiological distinction of action words in the fronto-central cortex. *Human Brain Mapping*, 21, 191–201.

Hutchinson, S., & Louwerse, M. (2013). Language statistics and individual differences in processing primary metaphors. *Cognitive Linguistics*, 24, 667–687.

Hutchinson, S., & Louwerse, M. M. (2014). Language statistics explain the spatial–numerical association of response codes. *Psychonomic Bulletin & Review*, 21, 470–478.

Liu, P., Connell, L., & Lynott, D. (2018). Can illness be bright? Linguistic distributional frequency and ease of simulation affect metaphor comprehension. Manuscript submitted for publication.

Louwerse, M. M. (2011). Symbol interdependency in symbolic and embodied cognition. *Topics in Cognitive Science*, 3, 273–302. http://dx.doi.org/10.1111/tops.2011.3.issue-2

Louwerse, M. M., & Connell, L. (2011). A taste of words: Linguistic context and perceptual simulation predict the modality of words. *Cognitive Science*, 35, 381–398.

Louwerse, M., & Hutchinson, S. (2012). Neurological evidence linguistic processes precede perceptual simulation in conceptual processing. *Frontiers in Psychology*, 385(3), 1–11.

Louwerse, M. M., & Jeuniaux, P. (2008). Language comprehension is both embodied and symbolic. In M. de Vega, A. Glenberg, & A. C. Graesser (Eds.), *Embodiment and meaning: A debate* (pp. 309–326). Oxford: Oxford University Press.

Louwerse, M. M., & Jeuniaux, P. (2010). The linguistic and embodied nature of conceptual processing. *Cognition*, 114, 96–104.

Lupyan, G., & Ward, E. J. (2013). Language can boost otherwise unseen objects into visual awareness. *Proceedings of the National Academy of Sciences*, 110, 14196–14201.

Lynott, D., & Connell, L. (2009). Modality exclusivity norms for 423 object properties. *Behavior Research Methods*, 41, 558–564.

Lynott, D., & Connell, L. (2010). Embodied conceptual combination. *Frontiers in Psychology*, *212*(1), 1–14.

Lynott, D., & Connell, L. (2013). Modality exclusivity norms for 400 nouns: The relationship between perceptual experience and surface word form. *Behavior Research Methods*, *45*, 516–526.

Molinaro, N., Barraza, P., & Carreiras, M. (2013). Long-range neural synchronization supports fast and efficient reading: EEG correlates of processing expected words in sentences. *Neurolmage*, *72*, 120–132.

Niedenthal, P. M. (2007). Embodying emotion. *Science*, *316*, 1002–1005.

Niedenthal, P. M., Barsalou, L. W., Winkielman, P., Krauth-Gruber, S., & Ric, F. (2005). Embodiment in attitudes, social perception, and emotion. *Personality and Social Psychology Review*, *9*, 184–211.

Riordan, B., & Jones, M. N. (2011). Redundancy in perceptual and linguistic experience: Comparing feature-based and distributional models of semantic representation. *Topics in Cognitive Science*, *3*, 303–345.

Saffran, J. R. (2003). Statistical language learning mechanisms and constraints. *Current Directions in Psychological Science*, *12*, 110–114.

Sanford, A. J., & Sturt, P. (2002). Depth of processing in language comprehension: Not noticing the evidence. *Trends in Cognitive Sciences*, *6*, 382–386.

Santos, A., Chaigneau, S. E., Simmons, W. K., & Barsalou, L. W. (2011). Property generation reflects word association and situated simulation. *Language and Cognition*, *3*, 83–119.

Simmons, W. K., Hamann, S. B., Harenski, C. N., Hu, X. P., & Barsalou, L. W. (2008). fMRI evidence for word association and situated simulation in conceptual processing. *Journal of Physiology-Paris*, *102*, 106–119.

Smith, L., & Yu, C. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition*, *106*, 1558–1568.

Solomon, K. O., & Barsalou, L. W. (2004). Perceptual simulation in property verification. *Memory & Cognition*, *32*, 244–259.

Vigliocco, G., Meteyard, L., Andrews, M., & Kousta, S. (2009). Toward a theory of semantic representation. *Language and Cognition*, *1*, 219–247.

Whorf, B. L. (1956). *Language, thought, and reality: Selected writings of Benjamin Lee Whorf*. Cambridge, MA: MIT Press.

Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, *9*, 625–636.

Winawer, J., Witthoft, N., Frank, M. C., Wu, L., Wade, A. R., & Boroditsky, L. (2007). Russian blues reveal effects of language on color discrimination. *Proceedings of the National Academy of Sciences*, 104, 7780–7785.

Xu, F. (2002). The role of language in acquiring object kind concepts in infancy. *Cognition*, *85*, 223–250.

Yee, E., Chrysikou, E. G., Hoffman, E., & Thompson-Schill, S. L. (2013). Manual experience shapes object representations. *Psychological Science*, *24*, 909–919.

Zwaan, R. A., & Taylor, L. J. (2006). Seeing, acting, understanding: Motor resonance in language comprehension. *Journal of Experimental Psychology: General*, *135*, 1–11.