Original Articles

Contextual predictability shapes signal autonomy

James Winters[a,*], Simon Kirby[b], Kenny Smith[b]

[a] *The Minds and Traditions Research Group, Max Plank Institute for the Science of Human History, Jena, Germany*
[b] *Centre for Language Evolution, University of Edinburgh, Edinburgh, United Kingdom*

ABSTRACT

Aligning on a shared system of communication requires senders and receivers reach a balance between *simplicity*, where there is a pressure for compressed representations, and *informativeness*, where there is a pressure to be communicatively functional. We investigate the extent to which these two pressures are governed by *contextual predictability*: the amount of contextual information that a sender can estimate, and therefore exploit, in conveying their intended meaning. In particular, we test the claim that contextual predictability is causally related to *signal autonomy*: the degree to which a signal can be interpreted in isolation, without recourse to contextual information. Using an asymmetric communication game, where senders and receivers are assigned fixed roles, we manipulate two aspects of the referential context: (i) whether or not a sender shares access to the immediate contextual information used by the receiver in interpreting their utterance; (ii) the extent to which the relevant solution in the immediate referential context is generalisable to the aggregate set of contexts. Our results demonstrate that contextual predictability shapes the degree of signal autonomy: when the context is highly predictable (i.e., the sender has access to the context in which their utterances will be interpreted, and the semantic dimension which discriminates between meanings in context is consistent across communicative episodes), languages develop which rely heavily on the context to reduce uncertainty about the intended meaning. When the context is less predictable, senders favour systems composed of autonomous signals, where all potentially relevant semantic dimensions are explicitly encoded. Taken together, these results suggest that our pragmatic faculty, and how it integrates information from the context in reducing uncertainty, plays a central role in shaping language structure.

1. Introduction

Reducing uncertainty about the intended meaning is fundamental to any good communication system (Piantadosi, Tily, & Gibson, 2012; Ramscar & Port, 2015). In achieving this aim, speakers and hearers need to coordinate with one another, relying not only on the creation of conventional forms, but also on the way these forms interact with the contextual information at hand (Clark, 1996; Croft, 2000; Lewis, 1969; Scott-Phillips, 2015; Sperber & Wilson, 1986). Without context, linguistic systems such as English would be woefully ambiguous, leaving the sentence *She passed the mole* uninterpretable as to whether the verb *passed* refers to *a form of motion* or *an act of giving* and whether the noun *mole* refers to *a small burrowing mammal*, *a person engaged in espionage*, *a Mexican sauce* or *a type of causeway*. In short, when the context is known and informative, it helps in reducing uncertainty (Piantadosi et al., 2012).

Context, in this sense, refers to the mutual cognitive environment (Sperber & Wilson, 1986) in which an utterance is situated and determines what is and is not informative for reducing uncertainty. This

consists of a figure (the target of interpretation), a ground (the immediate information brought to the act of interpretation), and a background (prior knowledge derived from previous frames) (Duranti & Goodwin, 1992; Terkourafi, 2009). And, as with any environment, the context will vary: some contexts are regular and predictable, whereas others fluctuate and are unpredictable. When viewed in this way, the context is a variable that determines the extent to which a speaker can estimate, and therefore exploit, information that is relevant to reducing uncertainty about their intended meaning – its *contextual predictability*.

For instance, if a speaker is providing directions to the nearest grocery store, then the context includes information in the immediate environment, such as the general direction of the store relative to the present position of the interlocutors, as well as background knowledge about how a hearer is likely to interpret an utterance given the outcomes of previous interactions. Predictable contexts are therefore those where the speaker is able to use information provided by the context to reduce uncertainty about their intended meaning: if the grocery store is near a park, and the speaker and hearer share knowledge about where this park is located, then saying "there's a grocery store about five

minutes away, next to the park where we play rugby" is sufficient for the hearer to find the grocery store. This is in contrast to a situation where the speaker and hearer are strangers and uncertainty exists as to the knowledge they both share with one another (e.g., the hearer is a tourist and does not know about the existence of a nearby park).

This relationship between context, meaning and uncertainty leads to an interesting trade-off in how linguistic systems are organised. Languages vary in their degree of *signal autonomy*: "the capacity for an utterance to be interpreted in isolation, without recourse to implicit linguistic, cultural, contextual or cotextual knowledge. Non-autonomous expression combines linguistic signals with context, pragmatics, paralinguistic signals and the like" (Wray & Grace, 2007: 556). One hypothesis is that autonomy is favoured in situations where speakers and hearers cannot rely on context for disambiguation (Kay, 1977): autonomous signals are advantageous insomuch as they reduce reliance on shared social, historical and local contexts in favour of internal structure (Hurford, 2011; Snow, Barnes, Chandler, Goodman, & Hemphill, 1991).

In this paper we present experimental evidence demonstrating that the degree of signal autonomy is causally related to contextual predictability: in an experiment where participants interact using an artificial language, highly predictable contexts favour systems composed of non-autonomous, context-dependent signals, whereas decreasing contextual predictability results in increased autonomy (context-independence). Crucially, we argue that these systems arise from the pressures of *informativeness* and *simplicity* (Kirby, Tamariz, Cornish, & Smith, 2015; Regier, Kemp, & Kay, 2015), with the degree of contextual predictability interacting with these two pressures to restrict the space languages explore.

1.1. Signal autonomy and contextual predictability

No natural language has completely autonomous signals in the sense of unambiguous clarity; context is always involved in reducing uncertainty about the intended meaning. But it is relatively uncontroversial to say there are degrees of autonomy. Contrast the use of indexical (context-dependent) and non-indexical (autonomous) forms of language: when referring to the day after today, English users will tend to say *tomorrow*, rather than the more autonomous counterpart of a specific date (e.g., *July 5th 2016*) (Hurford, 2011). Both are perfectly valid forms of expressing the relevant meaning, yet indexical forms require enrichment from external information (e.g., James lives on *this* street), whereas non-indexical forms are useful in providing specific information in the absence of such contexts (e.g., James lives on Milton Street).

It is not just individual constructions which vary in signal autonomy; languages, as collections of constructions, also vary in the extent to which they can be characterised as more or less autonomous (Hurford, 2011; Kay, 1977; Wray & Grace, 2007). An extreme example of this cross-linguistic variation in autonomy is found in Riau Indonesian – a colloquial variety of Malay/Indonesian with minimal syntactic structure and highly context-dependent expressions (for review see Gil, 2005). For instance, the combination of *ayam* ("chicken") and *makan* ("eat") yields a vast number of possible interpretations, ranging from *the chicken is eating* to *the chickens are eating* or *someone is eating the chicken* or even *someone is eating with the chicken* (Gil, 2005; Hurford, 2011). In short, the phrase *ayam makan* or *makan ayam* involves anything to do with chicken and eating; contextual information and inference do the rest of the work in sifting through possible interpretations.

1.2. The immediate context, the amount of shared context, and the historical context

Our general hypothesis is that variation in signal autonomy is predicted by the amount of contextual predictability. When the context is

Fig. 1. In context A, an English speaker can discriminate between both objects by using the cup or the bowl, whereas in context B they must use more elaborate expressions: *the metal cup* and *the wooden cup* (assuming the speaker obeys the rules of English for adjective use).

predictable, signal autonomy is low; when the context is unpredictable, signal autonomy is high. A complicating factor is determining what aspects of the context shape the degree of contextual predictability. This is problematic because operationalising separate aspects of context, and investigating the relationships between these aspects in a systematic fashion, is no simple task (see Bazire & Brézillon, 2005; Clark & Carlson, 1981). For our purposes, we restrict our focus to three types of context: *the immediate context, the amount of shared context*, and *the historical context*.

The immediate context refers to the situational information that is relevant for producing and comprehending an utterance. Consider the possible use of referring expressions in Fig. 1. Describing the object on the left in contexts A and B could be achieved with the referring expression *the metal cup* – this expression is capable of discriminating between referents in both contexts. Yet, based on a long history of psycholinguistic studies, it is only in context B where the expression *the metal cup* is used, with *the cup* being preferred when the adjective is not needed for discrimination (Olson, 1970; Pechmann, 1989; Sedivy, 2005; for review, see: Konopka & Brown-Schmidt, 2014). In this sense, a maximally predictable context is one in which a single semantic dimension (e.g., shape) is relevant for discrimination, with predictability decreasing as more dimensions are necessary for success (e.g., shape and material).

If the immediate context is the only relevant factor in determining contextual predictability, signal autonomy should pattern with the amount of contextual information in the local, situational context. In this case, low signal autonomy is expected when the situational context backgrounds some information (e.g., material) and highlights other information (e.g., shape). This narrow conception of context runs into problems when accounting for linguistic phenomena such as overspecification (where redundant, non-contrastive information is incorporated; see: Tinits, Nölle, & Hartmann, 2017). For instance, unlike material and scalar adjectives, which tend to pattern with the immediate context, colour adjectives are often used even when they are uninformative for discrimination[1] (e.g., Arts, Maes, Noordman, & Jansen, 2011; Rubio-Fernández, 2016; Sedivy, 2005). Similarly, languages often morphologically overspecify, obligatorily encoding

---

[1] A growing body of work into these *Redundant Colour Adjectives* (RCA) provides two explanations (Rubio-Fernández, 2016). First, the use of RCAs tends to be contingent on the semantic category, as evident in their presence for atypical objects (e.g., *the brown banana*) and clothes (e.g., collocations such as *black tie*) and their absence in typical (e.g., *the banana*) and geometrical figures (Dale & Reiter, 1995; Grodner & Sedivy, 2011). Second, speakers tend to produce RCAs when colour helps facilitate object recognition (e.g., polychrome versus monochrome displays), as well as when the language uses prenominal (e.g., English) as opposed to post-nominal adjectives (e.g., Spanish) (Rubio-Fernández, 2016).