Connecting content and logical words*

Emmanuel Chemla, Brian Buccola, Isabelle Dautriche

Abstract

Content words — e.g., nouns and adjectives — are generally connected: there are no gaps in their denotations; no noun means 'table or shoe' or 'animal or house'. We explore a formulation of connectedness which is applicable to content and logical words alike, and which compares well with the classic notion of monotonicity for quantifiers. On a first inspection, logical words satisfy this generalized version of the connectedness property at least as well as content words do — that is, both in terms of lexicons (with a study very narrowly restricted to English, however) and in terms of acquisition biases (with an artificial rule learning experiment). This reduces the putative differences between content and logical words, as well as the associated challenges these differences would pose, e.g., for learners.

1 Current generalizations separate content and logical words

No noun in English means 'bottle or eagle', and no quantifier means 'less than 5 or more than 10'. Generalizations about what words are possible and what words are not have been studied in two separate traditions: one is concerned with so-called 'content' words, and it can be found across psychological and philosophical studies of language; the other tradition is at the root of formal semantics and is concerned with 'logical' words such as connectives and quantifiers. We will argue against a strict separation of this kind and submit that at least some constraints proposed for one domain may apply nicely to the other. But let us first review the situation for content words and logical words separately.

1.1 Content words (nouns and adjectives): the phenomenology of connectedness

We will call 'content words' those words whose denotations can be construed as sets of objects: the set of eagles, the set of black entities, etc. Content words are typically nouns (*eagle*) and adjectives (*black*). The denotation of a content word generally satisfies a connectedness constraint: if two objects are *blickets*, then any object in between these two objects is also a *blicket*. This can be formalized as in the definition in (1). This definition presupposes the existence of a ternary in-between relation over objects, which we notate by $[a \textcircled{b} c]$ ('*b* is between *a* and *c*'). The exact nature of this relation is a question we will come back to later on. (The definition in (1) may also seem to depart from the classical notion of connectedness. In the appendix we situate it more precisely in the landscape of possible formalizations of the notion.)

(1) **Connectedness.** A set S is *connected* iff (if and only if) for any objects a, b, and c, if $a \in S$, $c \in S$, and $[a \textcircled{b} c]$, then $b \in S$.

Peter Gärdenfors vividly pushed forward the thesis that the denotation of a content word is generally connected (e.g., [9], but see [10] for an extension to cases beyond content words in the

For help, discussions, and criticisms, many thanks to Paul Egré, Roni Katzir, Manuel Križ, Philippe Schlenker, Raj Singh, Benjamin Spector, and Ede Zimmermann. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 313610, from the ESRC under the Future Research Leaders scheme (ES/N017404/1), and was supported by ANR-10-IDEX-0001-02 PSL and ANR-10-LABX-0087 IEC.

strict sense above). There are abundant exceptions (for a most recent critical approach, we refer to [12]; see also [11]), but there is also compelling evidence in favor of the role of such a constraint in the lexicons of natural languages (see [3] for discussion about linguistic universals and what arguments can be advanced, above and beyond the exceptions a potential universal may show).

The phenomenology typically associated with connectedness reaches beyond what can be observed from the lexicons of the languages of the world. In trying to guess the meaning of a new word, one goes beyond observed examples, presumably guided by several expectations about what words are like. The assumption that words are connected could play a crucial part in this generalization process: people who learn that X and Y are *blickets* would tend to assume that objects in between X and Y are also blickets; this has been shown to be true for both adults and children (e.g., [28]).[1]

Despite counter-examples, connectedness thus seems to constrain the denotation of (many) actual nouns and of words learned in artificial settings, with the learning bias and the fossilization observed in the lexicons being naturally understood as two sides of the same coin.

1.2 Logical words (connectives and quantifiers): their own classic constraints

Words do not always refer to sets of objects. Instead, denotations quite generally are conceived of as functions. For instance, instead of construing the denotation of a content word w as a set of objects S, one may equivalently construe w's denotation as a function f of type $\langle e, t\rangle$ — namely, that function f from the domain of entities (or objects) to the domain of truth values ($\{0,1\}$) such that $f(x) = 1$ iff $x \in S$. This function is known as the characteristic function of S.[2]

With such a type-theoretic approach to meanings, one can identify connectives as words with denotations of type $\langle t, \langle t, \langle \ldots, t\rangle \ldots\rangle\rangle$. For instance, the denotation of *and* may be seen as taking two truth values and returning a truth value, but, parsing the arguments in a different way, it can be seen as being of type $\langle t, \langle t, t\rangle\rangle$, in that it takes a first truth value, and then returns a function taking a second truth value before returning its output truth value. Similarly, one can study quantifiers, logical words of higher types. More precisely, (quantificational) determiners, e.g., *every*, have denotations of type $\langle et, \langle et, t\rangle\rangle$, and (generalized) quantifiers, e.g., *nobody*, have denotations of type $\langle et, t\rangle$. In a part of the literature that grew more from the formal semantics tradition than from psychology (although see [24]), it has been asked why connectives and quantifiers found in natural languages only cover a small subset of all possible functions of these types (e.g., [1, 14, 15, 25]), why some properties seem to be almost universal among quantifiers (e.g., conservativity) and are possibly definitional of logicality (see the work of [4, 27] about permutation invariance), and what properties are linguistically important (e.g., in the seminal work of [7, 22] about the role of monotonicity in explaining the distribution of polarity items). Hence, a couple of important formal properties of quantifiers, and more precisely, of quantificational determiners, have been identified by semanticists, among which:[3]

[1] Interestingly, people seem to not follow this generalization under some particular circumstances — namely, if the given exemplars of blickets, X and Y, are too dissimilar, then objects in between may not be counted as further blickets. But it turns out that a closer inspection strengthens the case in favor of the connectedness constraint: even in these conditions, people do not infer a non-connected meaning for the word *blicket*. To avoid postulating a non-connected meaning, they instead infer that there is homophony, i.e. two words with the same phonological form, each of which satisfies the connectedness constraint ([5, 6]).

[2] Since sets and their characteristic functions are equivalent notions, we often talk interchangeably about them, as is common in the formal semantics literature.

[3] Katzir and Singh ([18], following [20]), take a slightly different perspective and argue that logical words are always

(2) **Conservativity.** A function f of type $\langle et, \langle et, t\rangle\rangle$ is *conservative* iff for all sets A and B, $f(A)(B) = f(A)(A \cap B)$.

(3) **Permutation invariance.** A function f of type $\langle et, \langle et, t\rangle\rangle$ is *permutation invariant* iff for every permutation π of the domain of entities, $f(A)(B) = f(\pi(A))(\pi(B))$, where $\pi(X) = \{\pi(x) : x \in X\}$.

(4) **Monotonicity.** Given orders over σ and τ domains, a function f of type $\langle \sigma, \tau\rangle$ (e.g., a quantifier of type $\langle et, t\rangle$) is:

a. *upward monotonic* iff $a \leq b$ entails $f(a) \leq f(b)$,
b. *downward monotonic* iff $a \leq b$ entails $f(a) \geq f(b)$,
c. *monotonic* iff it is upward or downward monotonic.

Just as for connectedness of function words, these properties have been scrutinized in natural language lexicons, as well as in artificial learning studies, with the goal of trying to see whether the properties observed in lexicons may derive from learning biases (see for instance [16] on conservativity). New studies also try to provide computational arguments showing that some properties could create *a priori* formal advantages in learning (cf. [17, 26]).

1.3 Summary of the situation

Generalizations about content and logical words have been stated separately, with the underlying assumption that the two types of words follow different generalizations. If so, one should justify that a learner can navigate *a priori* through different generalizations, applying different constraints to different types of words. Researchers have sought to exhibit independent simple and observable facts about function and content words, such as their relative frequency, that learners could use to separate logical words from content words *a priori* (see [13]). We propose to take another direction. We propose to look at these constraints from a more general perspective, and first ask whether they can be translated to apply to words in general, not only the ones they were designed for. Once this is done, we can ask for any given property whether it is specific to some types of words or whether it actually applies across the board. Specifically, we will show that the connectedness constraint, most studied in connection to content words, is in fact very well suited to quantifiers.

2 Connectedness for quantifiers

We show how to extend the notion of connectedness discussed most specifically for content words, and how it applies to quantifiers very naturally to offer a close relative to monotonicity.

2.1 Connectedness of $\langle e, t\rangle$ functions, and beyond

As already mentioned, the denotation of a content word may be construed either as a set of entities or, equivalently, as the characteristic function of that set. Connectedness is naturally thought of as a property of sets. But it is of course equally possible to phrase it as a property of functions, and we propose to do so with the goal of applying this definition to words more generally.

---

obtained as combinations of two primitives (*min* and *max*). In this approach, except for the primitive elements, a property all words satisfy should be derived as a mere consequence of the properties of the primitives and of the combinatorial possibilities.

(5) **Connectedness (for $\langle e, t \rangle$ functions).** A function f of type $\langle e, t \rangle$ is *connected* iff for any objects a, b, and c, if $[a \textcircled{b} c]$, then $f(b) \geq f(a)$ or $f(b) \geq f(c)$.

We classically construe the domain of truth values (type t) as consisting of 1 (true) and 0 (false), with $0 < 1$. The connectedness of a set then corresponds to the connectedness of its characteristic function, as shown in (6). The definitions of connectedness thus coincide for the view of denotations of content words as sets or as functions, and one may thus talk about the connectedness of a content word directly without a risk of confusion.

(6) A set S of e objects is connected iff its characteristic function of type $\langle e, t \rangle$ is connected.

Proof. ($\rightarrow$) Let S be a connected set of objects and f its characteristic function. Let a, b, and c be any objects such that $[a \textcircled{b} c]$. If $a \notin S$ or $c \notin S$, then $f(a) = 0$ or $f(c) = 0$, hence $f(b) \geq f(a)$ or $f(b) \geq f(c)$. If $a \in S$ and $c \in S$, then $[a \textcircled{b} c]$ implies $b \in S$, too, hence $f(a) = f(b) = f(c) = 1$. Thus, f is connected. ($\leftarrow$) Let f be the connected characteristic function of a set of objects S. Let a, b, and c be any objects such that $a \in S$, $c \in S$, and $[a \textcircled{b} c]$. Then $f(a) = f(c) = 1$. Since $f(b) \geq f(a)$ or $f(b) \geq f(c)$, it follows that $f(b) = 1$, hence $b \in S$. Thus, S is connected. $\square$

This result is not specific to input domains of type e, and the proof extends straightforwardly to any input domain for which the notion of connectedness is defined — that is, which has an in-between relation for its input set and an ordering relation for its output set. The following definition extends the previous notion of connectedness to any function of type $\langle \sigma, \tau \rangle$, if the types come equipped with these requirements:[4]

(7) **Connectedness (general version).** A function f of type $\langle \sigma, \tau \rangle$ (with a ternary in-between relation $[\odot]$ over σ elements and an order $\leq$ over τ elements) is *connected* iff for any objects a, b, and c of type σ, if $[a \textcircled{b} c]$, then $f(b) \geq f(a)$ or $f(b) \geq f(c)$.

We may thus investigate connectedness for content words, but also for any other types of words, as long as the necessary in-between and order relations are available.

2.2 Connectedness for quantifiers: a comparison with monotonicity

The definition of connectedness is thus applicable to any word of type $\langle \sigma, \tau \rangle$ such that the input type σ has an in-between relation, and the output type τ has an ordering relation. The proper definition of an in-between relation is not in general an easy matter, and particularly so for objects: is a chair in between a sofa and a stool? Is an eagle in between a plane and a cat? Interestingly, however, an in-between relation is canonically given for quantifiers of type $\langle et, t \rangle$ because the natural subset relation over sets (or the equivalent relation over type $\langle e, t \rangle$ functions), $\subset$, induces an in-between relation. More generally, any strict partial order induces an in-between relation, as follows:

(8) **Canonical in-between relation induced by an order relation.** Let $<$ be a strict partial order over a set S. An in-between relation over S can be defined as: for any elements a, b, and c of S, b is in between a and c, notated as $[a \textcircled{b} c]_<$, iff $a < b < c$ or $c < b < a$.

[4]Other semantic properties can be extended to types beyond the cases they were initially thought for; see for instance [21] for an extension of the notion of homogeneity from $\langle e, t \rangle$ predicates across the type hierarchy.

Intuitively, a ternary relation constructed in this way from an order indeed should count as an in-between relation. More formally, any relation obtained in this way satisfies the axioms for betweenness proposed by [9] (see Appendix 1 for a proof).[5]

Thus, quantifiers have natural order relations for both their input and their output types, which induce a fully specified notion of connectedness. The general definition of connectedness in (7) applied to quantifiers (using the subset relation over the input set and the $0 < 1$ relation over the output set) is provided in (9).

(9) **Connectedness for quantifiers ($\langle et, t \rangle$ functions).** A function f of type $\langle et, t \rangle$ is *connected* iff for any sets of entities A, B, and C, if $[A \textcircled{B} C]_{\subset}$, $f(A) = 1$, and $f(C) = 1$, then $f(B) = 1$.

Because there is a natural order relation among sets, $\subset$ (from which we canonically derive an in-between relation), the more usual notion of monotonicity can also be defined for quantifiers (see (4)). Theorem 1 states the close proximity between the canonically derived notion of connectedness and the canonical notion of monotonicity.

(10) **Lemma 1.** If a quantifier is monotone, then it is connected.

Proof. Let Q be a monotone quantifier, and let A, B, and C be any sets such that $[A \textcircled{B} C]_{\subset}$, $Q(A) = 1$, and $Q(C) = 1$. Since $[A \textcircled{B} C]_{\subset}$, we have $A \subset B \subset C$ or $C \subset B \subset A$. If Q is upward monotone, then $A \subset B$ and $C \subset B$ each entail that $Q(B) = 1$, since both $Q(A)$ and $Q(C)$ equal 1; and if Q is downward monotone, then $B \subset C$ and $B \subset A$ each entail that $Q(B) = 1$, since both $Q(C)$ and $Q(A)$ equal 1. In all cases, whether Q is upward or downward monotone, and whether $A \subset B \subset C$ or $C \subset B \subset A$, we have $Q(B) = 1$. Thus, Q is connected. $\square$

(11) **Lemma 2.** If a quantifier is monotone, then its negation is connected.

Proof. Let Q be a monotone quantifier. Then the negation of Q is also monotone (albeit of opposite polarity), hence by Lemma 1, is connected. $\square$

(12) **Theorem 1.** A quantifier is monotone iff it is connected and its negation is connected.

Proof. ($\rightarrow$) Follows from Lemmas 1 and 2. ($\leftarrow$) Let Q be a connected quantifier whose negation, $\neg Q$, is also connected. Suppose one can find $A \subset B$ with $Q(A) = 1$ and $Q(B) = 0$, i.e. Q is not upward monotone. If $A = \varnothing$, then $Q(\varnothing) = 1$. Otherwise, $\varnothing \subset A \subset B$, which still entails that $Q(\varnothing) = 1$, for otherwise $\neg Q(\varnothing) = 1$ and $\neg Q(B) = 1$, which would entail that $\neg Q(A)$ is also 1, by connectedness of $\neg Q$, yielding a contradiction ($Q(A) = 1$ and $\neg Q(A) = 1$). Now, we can show that Q is downward monotone. Pick any A' and B' such that $A' \subseteq B'$ and $Q(B') = 1$. If $A' = B'$, or if $A' = \varnothing$, then $Q(A') = 1$. Otherwise, $\varnothing \subset A' \subset B'$, with the two extremes $\varnothing$ and B' yielding 1. Hence, by connectedness of Q, $Q(A') = 1$. Overall, we have shown that if Q is not upward monotone, then it is downward monotone. Thus, in all cases, Q is monotone. $\square$

[5] In (8) we defined the in-between relation in terms of the strict order $<$, rather than the non-strict order $\leq$, because the former, but not the latter, yields an in-between relation that satisfies [9]'s axioms for betweenness, which presupposes that betweenness is irreflexive. We could also just as well have used the non-strict order $\leq$, which would yield an in-between relation that departs from the axioms (B0 and B3, in particular), but would nevertheless still yield the exact same result (Theorem 1).

Theorem 1 makes explicit a previously unnoticed similarity between connectedness — a property previously discussed for content words, but not for quantifiers — and monotonicity — a property that has shown its importance in formal semantics. It states that monotonicity is *the* minimal property that ensures both connectedness and stability of connectedness under negation. Connectedness (but not, in general, monotonicity) is also stable under conjunction: the conjunction of two connected quantifiers is connected. Therefore, if there were pressure for meanings obtained compositionally (from conjunction and negation) to be connected, then we might expect 'primitive' expressions to generally be monotone: the conjunction of monotone expressions and/or of their negations is connected. Of course, it is still possible to obtain non-connected meanings, e.g., by first conjoining an upward monotone expression with a downward monotone one (which yields a connected but non-monotone expression), and then negating this whole expression. By De Morgan's laws, this is equivalent to taking a disjunction, which itself is a very productive way to compositionally produce non-connected meanings. Surely, there are situations in which we want to express a non-connected meaning, and disjunction is one way that language allows us to do so.[6] However, if there is a preference for connected meanings for primitive expressions and for expressions of low compositional complexity (those obtained by negation of connected quantifiers and by conjunctions of connected quantifiers and/or of their negations), then, as Theorem 1 establishes, having monotonic primitive expressions is one way to achieve this.

3 Connectedness for quantifiers: phenomenology and how to test it

As we have seen, there is a canonical way to define connectedness for quantifiers, and the obtained property is closely related to monotonicity (Theorem 1). Here we discuss how to investigate the empirical footprint of connectedness for quantifiers, as it has been explored for content words.

3.1 Consequence 1: lexicon

As discussed above, if the connectedness constraint is real, then we should observe that quantifiers lexicalized in the languages of the world are, more often than not, connected. Here, we offer a very first look at English, in the hope that more systematic studies, and broadly cross-linguistic studies, can follow. To do so, we will use a loose definition of lexicalization. First, note that generalized quantifiers are rarely lexicalized in English strictly speaking — that is, they are rarely expressed with a single word or morpheme; usually, quantificational determiners are (e.g., *every*). Most typically, generalized quantifiers come in the form of a combination of a quantificational determiner and a noun (e.g., *every teacher*). Second, with sufficient effort (and words), virtually any generalized quantifier can be expressed. So, rather than asking whether quantifier words typically correspond to connected quantifiers, we may instead ask whether connected quantifiers are typically expressed in simpler ways than non-connected quantifiers. The notion of simplicity would need to be made precise, but for the time being, as another limitation of this investigation, let us see whether an intuitive complexity measure can be useful.

[6] One might wonder why stability of connectedness under negation and conjunction somehow takes 'priority' over stability of connectedness under disjunction, but the reason is simple: the only way to ensure that *A or B* is always connected whenever *A* and *B* are each connected is to require that at least one of *A* or *B* is the tautological/contradictory quantifier, which would make such a disjunction vacuous/useless. Put differently, there's no substantive (non-vacuous) property that would ensure stability of connectedness under disjunction, and this in turn makes disjunction the natural way that language has to express non-connectedness.

Let us then observe some quantifiers:

(13) Some connected quantifiers of English:
somebody, *everybody*, *nobody*, *some people*, *all people*, *5 people*, *at most/at least/less than/more than/exactly 6 people*, *most/half of the people*, *only John*, all the negations of the preceding cases, *between 5 and 10 people*, *some but not all people*, *between a third and two thirds of the people*.

(14) Some non-connected quantifiers of English:
less than 5 or more than 10 people; *(exactly) 2*, *(exactly) 7*, *or (exactly) 10 people*; *none or all people*; *an odd/even number of people*.

Overall, there seems to be an appreciable difference between the connected and the non-connected quantifiers. Compare for instance the connected quantifier *between 5 and 10* and its non-connected negation *less than 5 or more than 10*: the former form does seem simpler than the latter. There is no counterpart of the *between* phrase, some word *between** such that *Between* 5 and/or 10 people came* would mean that less than 5 or more than 10 people came (e.g., *Outside of 5 and 10 people came* is not fluent English). In fact, the phrase *between 5 and 10 people* may be morphologically minimal: one needs to specify the 5 and 10 anchors, and this sets a lower bound on the complexity of ways to express such a quantifier (if this particular one were to be lexicalized, one might then expect an arbitrary number of similar expressions to be lexicalized, so for each pair (m, n), there would be a word meaning *between m and n*). So, this complex meaning requires an expression with some minimal complexity, but the complex expression we see here is otherwise fairly minimal: it contains as little as possible beyond the necessary boundaries 5 and 10. In fact, even the simpler *5 to 10 people came* would be interpreted in the same (connected) way, and here the complexity really seems minimal.[7] And conversely, just like there is no *between** word, there is no *to** word such that *5 to* 10* would be the non-connected negation of *5 to 10*.

Let us move to the more challenging case of *some but not all*. Here, there is no arbitrary boundary to be mentioned, and so one may have expected to have a simple expression for such a connected meaning, maybe even a single word. In fact, it seems here that its non-connected negation, *none or all*, is expressed with even fewer words. Interestingly, however, there *is* arguably a simpler way to express the meaning 'some but not all'—namely, *only some*. In fact, linguists have even argued that scalar implicatures, which may be described as a first approximation as the use of an implicit *only*, is what makes it unnecessary for *some but not all* to be part of a lexicon that already contains *some* and *all* [15]. Quite generally, connected but non-monotone meanings are expressible in a form such as *A and not B*, with *A* and *B* being monotone expressions of the same polarity—that is, as *only A*, with *B* implicitly being the part that *only* excludes. As a result, a lexicon for a language that is geared towards being able to express connected expressions as a priority may contain mostly monotone expressions, which ensures stability of connectedness under negation (cf. the discussion at the end of §2), and may rely on *only* or implicatures (or in extreme cases on explicit compositional mechanisms, even if costly) to create connected but non-monotone expressions.

The examples *an odd/even number of people* also seem to provide a challenge to the enterprise, given that they are not so complex, non-connected quantifiers. Note, however, that they come from

[7] And even the reduced *20, 30 people came* would be interpreted in the same way. This is presumably because of an implicit disjunction there, so it is similar to *20 or 30 people came*, and several views are possible about this, which may or may not make this relevant for the current purposes.

non-connected *content word* counterparts, the adjectives *odd* and *even*. Hence, one can hope that once the exceptions are understood for content words, they will also be justified in the quantifier case we are interested in.

This analysis remains preliminary, to say the least, certainly from a cross-linguistic perspective. A general inquiry of world languages would be very informative, but is beyond this study. We hope that others, in the spirit and with the help of resources like [19], might pursue them. At this point, we turn instead to the second side of the same coin: the immediate learnability of quantifiers that do and do not satisfy the connectedness constraint.

3.2 Consequence 2: artificial learning

Constraints on words may show in the lexicon of a language, but they may also be sought as biases in acquisition. Connectedness suggests a three-way distinction among quantifiers: a quantifier may be monotone (connected and its negation is connected), strictly connected (connected but its negation is not connected), or non-connected. Parallel to artificial learning studies concerned with connectedness for content words, we present a study that investigates whether this hierarchy among possible quantifiers surfaces as a bias in acquisition.

3.2.1 Task summary and hypothesis

The aim of the task is to learn a rule, which is quantificational in nature. Participants observed a series of displays containing a collection of five circles of different colors (see Figure 1). The rule relies on a quantifier that can be either monotone (e.g., "There are 3, 4, or 5 red objects"), connected (e.g., "There are 2, 3, or 4 red objects"), or non-connected (e.g., "There are 1, 2, or 4 red objects"). For each display, participants indicated whether the display is consistent with the rule (a yes/no response) and received feedback. We measured the number of trials needed to learn the rule and tested whether adults learn rules based on monotonic quantifiers the fastest, connected ones next, and non-connected ones the slowest.

3.2.2 Method

The data and the script for their analysis are available here: `http://semanticsarchive.net/Archive/WVhYzUwM/Chemla-Buccola-Dautriche-ConnectWords.html`.

Participants. 63 adults were recruited through Amazon's Mechanical Turk (25 females; $M = 34$ years; all native speakers of English) and compensated $1.80 for their participation. All participants were included in the analysis and randomly assigned to one of nine possible rules ($N = 9$ per rule); see Conditions below.

Procedure and display. Participants were tested online. They were instructed that they were to learn a rule by being exposed to a series of trials containing a collection of objects. For each trial, they would have to decide whether the display is consistent with the rule by pressing a "yes" or a "no" button and would receive feedback on their response. In the instructions, participants were shown an example of a trial. As represented in Figure 1, each trial consisted of a collection of 5 circles aligned horizontally on a white background. The circles could be one of 5 different colors. The feedback was displayed in a horizontal bar positioned at the top of the screen. The feedback bar turned green and displayed the prompt "Correct!" for correct responses or turned red and displayed the word "Incorrect" for incorrect responses. Correct responses received 2 seconds of feedback before the next trial commenced. Incorrect responses were penalized with 6 seconds of feedback to increase attention to the task.

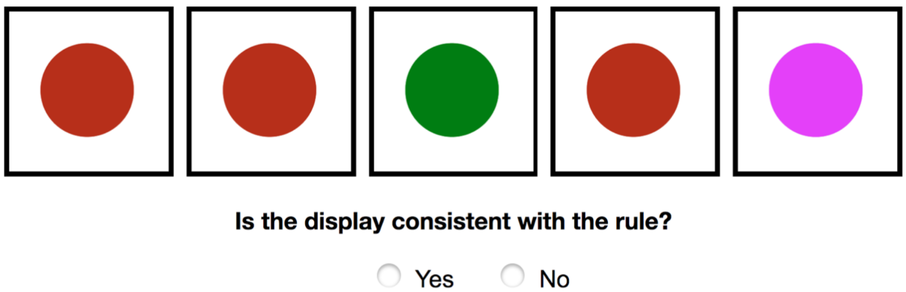

Figure 1. Example of a trial.

Conditions. Each participant learned a single rule. The rule could be either monotone, connected (but not monotone), or non-connected. For each of these conditions, there were two or three possible rules as listed in Table 1 below.

Condition	Rules
Monotone	"There are 0, 1, or 2 red circles."
	"There are 3, 4, or 5 red circles."
Connected	"There are 2, 3, or 4 red circles."
	"There are 1, 2, or 3 red circles."
Non-connected	"There are 0, 1, or 5 red circles."
	"There are 0, 4, or 5 red circles."
	"There are 1, 2, or 4 red circles."

Table 1. Participants learned one of five rules, belonging to one of three conditions: monotone, connected (but not monotone), or non-connected.

The monotone rules are true for displays containing any number of red circles in a set of 3 contiguous numbers (which guarantees connectedness), where the list includes either the lowest (0) or the highest (5) number of possible red circles (which guarantees monotonicity). The connected rules are true for 3 contiguous numbers of red circles on the screen, but they do not include the lowest or the highest boundaries of possible red circles, so that it is not a monotone rule. The non-connected rules are true for 3 non-contiguous numbers of red circles. The first two rules are the negation of the connected rules, and the third one is an additional case added where both the rule and its negation are non-connected.

Design and learning criteria. Participants learned in blocks of 6 trials, each displaying a different number of red circles (0 to 5, in random order). For each block, there were the same number of correct "yes" and "no" responses, independently of the rule being learned (all contained 3 values for true, 3 values for false). Participants were exposed to a minimum of 3 blocks (18 trials). The experiment ended when participants reached the learning criteria: either by answering correctly to all trials of a block, or after having answered a total of 250 trials ($N = 1$). After the experiment, they filled out a small questionnaire asking whether they could formulate the rule explicitly and, if so, to do so in their own words.

Data analysis. To quantify the ease with which different rules are learned, we fit the number of trials needed to learn the rule using a mixed model in R (`lme4` package [2]). The model included a categorical predictor Connectedness (monotone, connected, non-connected) as well as a random intercept for each rule. The model was specified as: `Ntrials ~ Connectedness + (1 | Rule)` and compared to a model without the predictor Connectedness to establish its effect on learning biases.

3.2.3 Results and analysis

Figure 2 reports the average number of trials needed to learn a rule per Connectedness (monotone, connected, and non-connected).

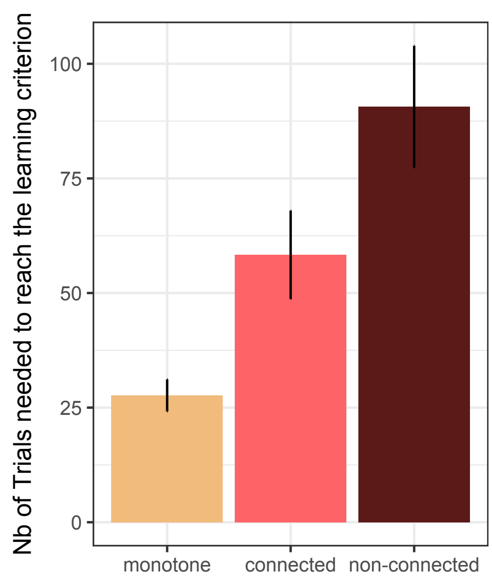

Figure 2. Average number of trials needed to reach the learning criterion for each Connectedness condition (monotone, connected, non-connected). Error bars indicate standard errors of the mean.

Analysis 1. On average, participants learned monotone rules the fastest ($M = 28$ trials; $SE = 4$), connected rules next ($M = 58$ trials; $SE = 10$), and non-connected rules the slowest ($M = 91$ trials; $SE = 13$), following the predicted acquisition biases. Yet, Connectedness was only a marginal predictor of learning performance ($\chi^2 = 5.26$; $p = 0.07$). Inspection of the learning performance per individual rules (see Table 1) revealed differences in participants' learning speed in the non-connected condition: both learning performances for rules (a) "There are 0, 1, or 5 red circles" and (b) "There are 0, 4, or 5 red circles" were comparable to the connected rules ($M = 64$ trials; $SE = 15$), but learning the rule (c) "There are 1, 2, or 4 red circles" took considerably longer ($M = 144$ trials; $SE = 27$). As a result, most of the difference between conditions is captured by the random intercept for each rule in our statistical model.

Analysis 2 (post-hoc). Note, however, that the negation of rules (a) and (b) are both connected (and these negations are the two connected rules that participants learned in the connected condition), while the negation of the harder rule (c) is not. In a sense, then, rules (a) and (b) are more connected than rule (c). Furthermore, in a binary yes/no task like the present one, it makes sense for a rule and its negation to be learned in a similar way, the difference between a rule and its negation being simply a matter of mapping the yes/no buttons in an opposite manner. A post-hoc analysis, then, in which the Connectedness factor codes whether neither, one, or both the rule and its negation are connected, may be pursued, and it yields a significant effect: $\chi^2 = 19.1$; $p < .001$.[8]

Analysis 3 (dynamic). We may also pursue a different test of the role of connectedness, one that would be more independent of the actual rules tested in the experiment and that looks more at the dynamic of learning in time. We looked at whether participants' decision to apply a rule when n red circles are present on the screen is dependent on their responses when $n-1$ or $n+1$ red circles are present within the same block of trials. Participants were more likely to apply the rule when n red circles were present if they applied it to both configurations where $n-1$ or $n+1$ red circles were present ($M = 0.68$; $SE = 0.02$) than if they applied it to only one of the surrounding configuration (i.e., $n-1$ or $n+1$; $M = 0.54$; $SE = 0.02$; $\beta - 0.51; z = -5.14; p < .001$) or none ($M = 0.40$; $SE = 0.02$; $\beta - 0.98; z = -8.10; p < .001$), and this even when controlling for the connectedness of the rule being tested.[9] This analysis thus directly reveals a bias for connectedness during learning, independently of the actual rule to be learned.

3.2.4 Discussion and conclusion

Participants learned rules based on monotonic quantifiers the fastest, connected ones next, and non-connected ones the slowest. Interestingly, they also show a general bias to postulate connected meanings during learning, and they did so independently of the actual rule to be learned.

Is the task done linguistically? It could be the case that learners find it easier to learn monotonic and connected rules over non-connected ones simply because they could formulate easier linguistic expressions to remember these rules. For instance, the monotonic rule "There are 0, 1, or 2 red

[8]We note that rules (a) and (b) differ from rule (c) in another way — the gap is smaller in rule (c). As a result, they may trigger different learning mechanisms and representations, whereby one is learned as a connected rule plus some exception somewhere further away, while the other is learned as a genuinely non-connected rule. This would be coherent with the work on biases for homophones discussed in footnote 1. We do not explore this possibility further here, and favor a more direct look at connectedness with Analysis 3.

[9]We modeled participants' choice in applying the rule when n red circles where present (yes or no) using a mixed logit model specified as `choiceN ~ NContiguousResponses + Connectedness + (1 | Subject) + (1 | Rule)`. The predictor NContiguousResponses is the sum of "yes" responses (0 to 2) for both trials where $n+1$ and $n-1$ circles were present within the same block.

circles" could be expressed as *There are at most 2 red circles*, and the connected rule "There are 2, 3, or 4 red circles" as *There are between 2 and 4 red circles*; but for the non-connected rules, there is no choice for participants but to remember a list of 3 non-contiguous numbers. We find this possibility unlikely as only one third of the participants could provide an explicit formulation of the rule they learned. It is also worth mentioning that we observed differences between rules for which there is no obvious asymmetry between the different English expressions needed to express them (the optimal ways to express the connected and non-connected rules that gave rise to different speed of acquisition seem comparable).

These findings thus suggest that connectedness, a natural formal constraint on word meanings, may be an active constraint during language acquisition, biasing the learning device to search for connected meanings first.

4 General discussion

Connectedness is typically defined for nouns. However, thanks to the natural order relation among sets, $\subset$, which induces a canonical in-between relation, the definition is actually more straightforward for quantifiers than for nouns, where one needs to converge first (or simultaneously) on the appropriate in-between relation among objects. Once defined, connectedness for quantifiers is very close to the usual notion of monotonicity, a notion that plays an important linguistic role. We also provided preliminary evidence that connectedness could enjoy the phenomenology expected for such constraints, namely simplicity of expressions that satisfy them, and biases in learning.

Monotonicity is defined not only for generalized quantifiers, but also for connectives and quantificational determiners. One would thus need to extend the inquiry about connectedness to these types. This comes with challenges. First, for connectives, the t type is so simple that any relevant order or in-between relation would not help yield very strong constraints. Second, for both connectives and determiners, one may seek a notion that could apply to the different arguments separately (as for monotonicity) or jointly (considering that the input could be a pair of arguments for that matter, even if it would not be so from a syntactic perspective). This requires a couple of choice points then, which would definitely be worth investigating in detail. In fact, in the domain of content words, extension of the inquiry to more types was one of the achievements of [10].

5 Conclusion

We offer to evaluate the generality of the following claim: all words are connected, whether they are logical or content words. We showed how to relate the connectedness property discussed for content words and the monotonicity property of quantifiers, the latter being a particular, stronger form of the former, which is most natural in the presence of a canonical ordering relation (Theorem 1).

More generally, we propose to search for universals about word meanings (see [8] for a review of universals in semantics). The grand goal is to find a list of properties which are, in some sense, double universals: universals across languages, but also across word types (as soon as the type meets the requirements for the property to be expressed, e.g., whenever some in-between relation is available). There will be misses, due to accidents of various sorts, but one ought to seek other evidence that a property is active, e.g., in language history ([3]) or in (actual and artificial) acquisition. Interesting predictions could be obtained about words with other types (e.g., adjectives which may be of types $\langle et, et \rangle$) or about how languages may lexicalize different 'quantifiers', if quantification is expressed in different ways and with different types ([23] is most relevant here).

There may exist an implicit assumption that logical and content words cannot be subject to similar universals.[10] We proposed a case study where we show that this assumption is inappropriate, showing in the case of connectedness (content words) and monotonicity (logical words) that it might lead us to miss a generalization concerning the similarity between content and logical words.

Appendix 1: Gärdenfors's axioms for betweenness

We prove that for any strict partial order $<$ (over a set S), the associated in-between relation $[\odot]_<$ (defined as $[a \textcircled{b} c]_<$ iff $a < b < c$ or $c < b < a$, see (8)) satisfies the four axioms of betweenness proposed by Gärdenfors in [9].

(15) **B0.** If $[a \textcircled{b} c]_<$, then a, b, and c are distinct objects.

Proof. Let a, b, and c be any objects such that $[a \textcircled{b} c]_<$. Then $a < b < c$ or $c < b < a$ and, because the order is strict, a, b, and c must be distinct objects. □

(16) **B1.** If $[a \textcircled{b} c]_<$, then $[c \textcircled{b} a]_<$.

Proof. For any a, b, and c, both $[a \textcircled{b} c]_<$ and $[c \textcircled{b} a]_<$ are defined as $a < b < c$ or $c < b < a$. □

(17) **B2.** If $[a \textcircled{b} c]_<$, then not $[b \textcircled{a} c]_<$.

Proof. Let a, b, and c be any objects such that $[a \textcircled{b} c]_<$. Then $a < b < c$ or $c < b < a$. If $a < b < c$, then it cannot be the case that $b < a < c$ (because $a < b$) or that $c < a < b$ (because $a < c$), i.e. it cannot be the case that $[b \textcircled{a} c]_<$. The reasoning is the same if $c < b < a$. □

(18) **B3.** If $[a \textcircled{b} c]_<$ and $[b \textcircled{c} d]_<$, then $[a \textcircled{b} d]_<$.

Proof. Let a, b, c, and d be any objects such that $[a \textcircled{b} c]_<$ and $[b \textcircled{c} d]_<$. Then $a < b < c$ or $c < b < a$, and $b < c < d$ or $d < c < b$. There are four cases to consider:

1. $a < b < c$ and $b < c < d$: This entails that $a < b < d$. Hence, $[a \textcircled{b} d]_<$.
2. $a < b < c$ and $d < c < b$: This is a contradiction (cf. the ordering constraints for b and c).
3. $c < b < a$ and $b < c < d$: This is a contradiction (cf. the ordering constraints for b and c).
4. $c < b < a$ and $d < c < b$: This entails that $d < b < a$. Hence, $[a \textcircled{b} d]_<$.

[10] One superficial reason for such an assumption is that one goal that is sometimes pursued is to propose a set of universals which together exclude any word that is not attested in a natural language. For such a goal, there surely is an asymmetry between logical words and content words, the latter being potentially infinite in number and productively created. But this asymmetry is artificial. It would be wrong to think that because such goals may apply to logical words more than to content words, the study of universals must be divorced at all for logical and content words.

In all non-contradictory cases (i.e. the only way for the premises to hold), $[a \textcircled{b} d]_<$. □

(19) **B4.** If $[a \textcircled{b} d]_<$ and $[b \textcircled{c} d]_<$, then $[a \textcircled{b} c]_<$.

Proof. Let a, b, c, and d be any objects such that $[a \textcircled{b} d]_<$ and $[b \textcircled{c} d]_<$. Then $a < b < d$ or $d < b < a$, and $b < c < d$ or $d < c < b$. There are four cases to consider:

1. $a < b < d$ and $b < c < d$: This entails that $a < b < c$. Hence, $[a \textcircled{b} c]_<$.
2. $a < b < d$ and $d < c < b$: This is a contradiction (cf. the ordering constraints for b and d).
3. $d < b < a$ and $b < c < d$: This is a contradiction (cf. the ordering constraints for b and d).
4. $d < b < a$ and $d < c < b$: This entails that $c < b < a$. Hence, $[a \textcircled{b} c]_<$.

In all non-contradictory cases (i.e. the only way for the premises to hold), $[a \textcircled{b} d]_<$. □

Appendix 2: in-between relations, paths, and various notions of connectedness

Instead of starting from an in-between relation, one could use paths. This will allow us to make the relation between different notions of connectedness more transparent. Under appropriate conditions on the underlying space S, there are several ways to define paths. Here are some:

(20) The path from a to b is the segment from a to b (assuming 'segments' can be defined).

(21) For any x, the x-path from a to b is the union of the segments from a to x and from b to x.

(22) A path from a to b is (the image of) a continuous function π from $[0, 1]$ to S, such that $\pi(0) = a$ and $\pi(1) = b$.

(23) The path from a to b is the set of elements x such that $a \leq x \leq b$ or $a \geq x \geq b$.

For any of these notions, P, let us note $P(a, b)$ the set of paths from a to b. One can define corresponding notions of P-connectedness of a set, as the existence of a P-path between any two elements. As before (result (6)), we can do this in two equivalent ways by looking at the set itself or at its characteristic function:

(24) A subset X of S is P-connected iff $\forall a, b \in X : \exists \pi \in P(a, b) : \pi \subseteq X$.

(25) A function f is P-connected iff $\forall a, b : \exists \pi \in P(a, b) : \forall x \in \pi : f(x) \geq f(a)$ or $f(x) \geq f(b)$.

Thus, a notion of path, P, gives rise to a notion of P-connectedness, which may correspond to a classic notion with a more specific name: it is convexity with definition (20) for paths, star-shapedness with definition (21), and usual connectedness with definition (22).

From a notion of path, one may derive an in-between relation. In particular when $P(a, c)$ always contains a single element, we can define $[a \textcircled{b} c]_P$ as b being in *the* path $P(a, c)$. When there may be more than one path in $P(a, c)$, there are several possibilities. For this reason, the notion of path given in (22), the one corresponding best to classic connectedness, is also the hardest to associate with an in-between relation because there may be multiple paths within a given $P(a, c)$. Let us not dive into these issues here. Most importantly, we note that in the main part of the text, through (8) in particular, we derive the in-between relation from an order, as in (23) (in (8) we use the corresponding strict order, but this is inconsequential). The corresponding notion of connectedness can now be located in the space of possible comparable notions.

References

[1] Jon Barwise and Robin Cooper. Generalized quantifiers and natural language. *Linguistics & Philosophy*, 4(2):159–219, 1981.

[2] Douglas Bates, Martin Maechler, Ben Bolker, Steven Walker, et al. lme4: Linear mixed-effects models using eigen and s4. *R package version*, 1(7):1–23, 2014.

[3] Balthasar Bickel. Linguistic diversity and universals. In Enfield, N J, Kockelman, Paul, and Sidnell, Jack, editors, *The Cambridge Handbook of Linguistic Anthropology*, pages 101–124. Cambridge: Cambridge University Press, 2014.

[4] Denis Bonnay. Logicality and invariance. *Bulletin of Symbolic Logic*, 14(1):29–68, 2008.

[5] Isabelle Dautriche and Emmanuel Chemla. What homophones say about words. *PloS one*, 11(9):e0162176, 2016.

[6] Isabelle Dautriche, Emmanuel Chemla, and Anne Christophe. Word learning: Homophony and the distribution of learning exemplars. *Language Learning and Development*, 12(3):231–251, 2016.

[7] Giles Fauconnier. Polarity and the scale principle. *Chicago Linguistics Society*, 11:188–199, 1975.

[8] Kai von Fintel and Lisa Matthewson. Universals in semantics. *The linguistic review*, 25(1-2):139–201, 2008.

[9] Peter Gärdenfors. *Conceptual spaces: The geometry of thought*. MIT press, 2004.

[10] Peter Gärdenfors. *The geometry of meaning: Semantics based on conceptual spaces*. MIT Press, 2014.

[11] Christopher Gauker. A critique of the similarity space theory of concepts. *Mind & Language*, 22(4):317–345, 2007.

[12] José V. Hernández-Conde. A case against convexity in conceptual spaces. *Synthese*, pages 1–27, 2016.

[13] Jean-Rémy Hochmann. Word frequency, function words and the second gavagai problem. *Cognition*, 128(1):13–25, 2013.

[14] Laurence Horn. A natural history of negation. *University of Chicago Press, Chicago*, 1989.

[15] Laurence Robert Horn. *On the semantic properties of logical operators in English*. PhD thesis, University of California, Los Angeles, 1972.

[16] Tim Hunter and Jeffrey Lidz. Conservativity and learnability of determiners. *Journal of Semantics*, 30(3):315–334, 2013.

[17] Roni Katzir and Noa Peled. Compression-based learning of natural-language quantifiers. Ms, 2017.

[18] Roni Katzir and Raj Singh. Constraints on the lexicalization of logical operators. *Linguistics & Philosophy*, 36(1):1–29, 2013.

[19] Edward Keenan and Denis Paperno. *Handbook of quantifiers in natural language*, volume 90. Springer, 2012.

[20] Edward L. Keenan and Jonathan Stavi. A semantic characterization of natural language determiners. *Linguistics & Philosophy*, 9:253–326, 1986.

[21] Manuel Križ. *Aspects of Homogeneity in the Semantics of Natural Language*. PhD thesis, University of Vienna, 2015.

[22] William A. Ladusaw. *Polarity Sensitivity as Inherent Scope Relations*. PhD thesis, University of Texas at Austin, Austin, TX, 1979.

[23] Lisa Matthewson. Quantification and the nature of crosslinguistic variation. *Natural Language Semantics*, 9(2):145–189, 2001.

[24] Steven T Piantadosi, Joshua B Tenenbaum, and Noah D Goodman. The logical primitives of thought: Empirical foundations for compositional cognitive models. *Psychological review*, 123(4):392, 2016.

[25] Pieter AM Seuren and Dany Jaspers. Logico-cognitive structure in the lexicon. *Language*, 90(3):607–643, 2014.

[26] Shane Steinert-Threlkeld and Jakub Szymanik. Learnability and semantic universals. Ms. University of Amsterdam, 2018.

[27] Alfred Tarski. What are logical notions? *History and Philosophy of Logic*, 7:143–154, 1986.

[28] Fei Xu and Joshua B Tenenbaum. Word learning as Bayesian inference. *Psychological review*, 114(2):245, 2007.