Cognition 181 (2018) 141–150

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Original Articles

Comprehenders model the nature of noise in the environment

Rachel Ryskin[a,b,*], Richard Futrell[c], Swathi Kiran[b], Edward Gibson[a]

[a] Dept. of Brain and Cognitive Sciences, Massachusetts Institute of Technology, 43 Vassar St., Cambridge, MA 02139, United States
[b] Dept. of Speech, Language, and Hearing Science, Boston University, 635 Commonwealth Avenue, Boston, MA 02215, United States
[c] Dept. of Language Science, University of California, Irvine, 3151 Social Sciences Plaza, Irvine, CA 92697, United States

ARTICLE INFO

Keywords:
Sentence comprehension
Noisy-channel
Rational inference
Adaptation
Error correction

ABSTRACT

In everyday communication, speakers make errors and produce language in a noisy environment. Recent work suggests that comprehenders possess cognitive mechanisms for dealing with noise in the linguistic signal: *a noisy-channel model*. A key parameter of these models is the *noise model*: the comprehender's implicit model of how noise affects utterances before they are perceived. Here we examine this noise model in detail, asking whether comprehension behavior reflects a noise model that is adapted to context. We asked readers to correct sentences if they noticed errors, and manipulated context by including exposure sentences containing obvious deletions (*A bystander was rescued by the fireman in the nick time.*), insertions, exchanges, mixed errors, or no errors. On test sentences (*The bat swung the player.*), participants' corrections differed depending on the exposure condition. The results demonstrate that participants model specific *types* of errors and make inferences about the intentions of the speaker accordingly.

1. Introduction

Everyday language use occurs amid myriad sources of noise. In a conversation, the speaker may say one word when she intended to say another, there may be other conversations going on in the same room, and the listener may mishear what was said. Each of these types of noise serves to corrupt the signal that is transmitted from speaker to listener (Shannon, 1948). One might think that such noise would pose major impediments to efficient communication. Yet language comprehension typically unfolds without noticeable effort.

Because of this noise, comprehenders maintain uncertainty about the nature of preceding words. When reading sentences such as, "The coach smiled at the player tossed the ball" readers' eye movements indicate that they leave open the possibility that "at" was actually "and." Replacing "at" with "and" allows the interpretation of "tossed" as a finite verb rather than a past participle; the former interpretation has a much higher conditional probability (Levy, Bicknell, Slattery, & Rayner, 2009). Thus, readers have probabilistic representations of language input—in particular, syntactic constructions—and use prior knowledge to infer the intended meaning.

Recent theories have proposed that the language processing system maintains uncertainty about the input because it is designed to optimally decode the intended meaning from a signal transmitted over a noisy channel (Bergen, Levy, & Gibson, 2012; Gibson, Bergen, & Piantadosi, 2013; Jaeger, 2010; Levy et al., 2009; Levy, 2008). In particular, Gibson et al. (2013) lay out a framework for sentence comprehension that entails the rational (Bayesian) integration of noisy evidence and semantic priors. On their account, the producer chooses an intended sentence s_i in order to communicate her intended meaning, m_i. s_i is conveyed across a noisy channel and is corrupted by noise originating from the producer, comprehender, or environment. The comprehender perceives sentence s_p and tries to infer s_i. Communication succeeds when the intended sentence s_i can be recovered from s_p. This process can be formalized by considering an ideal observer (Geisler & Diehl, 2003) model of language comprehension, where the comprehender engages in optimal Bayesian decoding of the intended meaning:

$$P(s_i|s_p) \propto P(s_i)P(s_i \rightarrow s_p) \quad (1)$$

In Eq. (1), $P(s_i|s_p)$ represents the probability assigned by the comprehender to any hypothesized s_i, given the observed linguistic input s_p. By Bayes rule, this probability can be rewritten as the prior probability $P(s_i)$ that a producer would wish to communicate s_i, multiplied by the probability of s_i being corrupted to s_p during communication, $P(s_i \rightarrow s_p)$. The prior, $P(s_i)$, represents the comprehender's relevant linguistic and world knowledge, and biases comprehenders towards a priori plausible utterances. The noise model $P(s_i \rightarrow s_p)$ encodes the comprehender's knowledge of how sentences can be

* Corresponding author at: Dept. of Brain and Cognitive Sciences, Massachusetts Institute of Technology, 43 Vassar St., Cambridge, MA 02139, United States.
E-mail address: ryskin@mit.edu (R. Ryskin).

https://doi.org/10.1016/j.cognition.2018.08.018
Received 12 July 2017; Received in revised form 24 August 2018; Accepted 28 August 2018

corrupted—for instance, smaller changes to a sentence are more likely than larger ones. By integrating P(s_i) and P($s_i \rightarrow s_p$), comprehenders may arrive at interpretations which differ from the literal meanings of the acoustic or visual stream. That is, if comprehenders perceive an implausible sentence s_p (e.g., The oven cleaned the grandmother) which is "close" to a more plausible sentence (e.g., The grandmother cleaned the oven), they should infer that the producer actually uttered (or intended to utter) the plausible version.

Gibson et al. (2013) provide evidence for several predictions of the noisy-channel framework in a series of experiments where participants read implausible sentences (e.g., The oven cleaned the grandmother) followed by comprehension questions (e.g., Was the grandmother cleaned by someone/something?), which probed whether participants interpreted the sentence literally (answer: Yes) or inferred that the intended sentence had been corrupted (answer: No). They found that comprehenders were (a) more willing to forego the literal interpretation when the semantically plausible interpretation involved positing fewer changes, (b) more likely to infer nonliteral meanings when the change involved a deletion compared to an insertion, consistent with the Bayesian size principle (Xu & Tenenbaum, 2007), (c) more likely to endorse literal interpretations when the fillers contained errors, indicating that they had inferred a higher noise rate; and (d) less likely to endorse literal interpretations when the base rate of implausible sentences was increased, suggesting that they had adjusted their semantic prior. Further, Poppels and Levy (2016) replicated these results and demonstrated that, in addition to deletions and insertions, word exchanges represent a likely form of corruption (e.g., The package fell from the floor to the table.).

1.1. Noise variation

Gibson et al. (2013) demonstrated that participants adapt their noise model when provided with evidence of a high base-rate of syntactic errors. Further, listeners infer a higher noise rate when listening to foreign-accented speech (Gibson et al., 2017). Yet, how the noise likelihood term ($s_i \rightarrow s_p$) responds to input characteristics beyond error rate has yet to be explored. Critically, we can ask: is the noise model sensitive to the nature of errors or simply to the rate of errors?

In real-world language use, many properties of the noise, beyond the rate, vary with context. For example, second language (L2) learners may make certain errors in English that a native speaker is unlikely to make and that are influenced by their native language (see MacWhinney, 1992). Native speakers of Russian tend to omit articles when speaking L2 English (e.g., Ionin, Ko, & Wexler, 2004), while native speakers of French may exchange the orders of adjectives and nouns in L2 English (Nicoladis, 2006). If the comprehender's noise model is sensitive to the nature of errors, it will have different properties when listening to an L2 English speaker from Russia than to an L2 English speaker from France. However, if the noise model is sensitive to an overall rate of errors, it will be similar for the two speakers.

Recent findings suggest that comprehenders rapidly learn and adapt to the linguistic patterns (e.g., frequencies of syntactic constructions, phonetic category boundaries) present in their environment in order to achieve more efficient language processing (Fine, Jaeger, Farmer, & Qian, 2013; Kleinschmidt & Jaeger, 2015; Ryskin, Qi, Duff, & Brown-Schmidt, 2017; though see Harrington Stack, James, & Watson, 2018 for an example of limits on this ability). Similarly, comprehenders may track the types of errors they perceive in a given environment and rapidly adapt the likelihoods of components of the noise model. For example, after hearing a speaker repeatedly drop articles (e.g., "We had nice time at beach."), the listener's noise model may put high probability on certain words being deleted, but the probability of insertions may not change. Thus, the noise model for the article-dropping speaker would have a larger ratio of deletions to other errors, as compared to the noise model for a generic, native English speaker. Forming these fine-grained, context-specific representations of the noise would likely allow comprehenders to make more accurate inferences about the intended meaning s_i. We call such a noise model a context-specific noise model.

On the other hand, hearing a speaker repeatedly drop articles may lead the listener's noise model to put higher probability on all possible errors, perhaps on the reasonable assumption that a speaker who makes one type of error is likely to also make other errors in the future. Under such a context-invariant noise account, the comprehender's noise model always possesses the same general properties (e.g., more edits are less likely than fewer edits, insertions are less likely than deletions) and varies only in the base-rate of corruptions. In an environment with a high base-rate of errors, comprehenders simply increase the likelihoods of all errors by a constant. In every other respect, the probabilities of different occurrences (e.g., deletions vs. insertions) maintain the same ratio across contexts. Inferring the noise model would then simply reflect the process of adjusting all the likelihoods in the noise model up or down, depending on recent evidence.

Whether comprehenders have context-invariant or context-specific noise models gets at the more general question of how people trade off complexity of models and accuracy in prediction. If the context-invariant model is correct, then this suggests that comprehenders weight model simplicity as more important than accuracy in prediction: the context-invariant noise model only has one parameter, the noise rate, and thus it should be easier to learn and deploy than a more complex model. If the context-specific model is true, then comprehenders weight accuracy as more important than model complexity in this case: the context-specific model achieves higher accuracy at the cost of greater complexity. The optimal tradeoff of accuracy and complexity will depend on the true rate of context-specificity in the world and on the exact nature of the complexity cost for noise models. These complexity-accuracy tradeoffs are at the heart of all theories of statistical learning (MacKay, 2003; Solomonoff, 1964). Investigating these two particular hypotheses in the context of noisy-channel language understanding allows us to develop models of how complexity and accuracy trade off in language processing.

In the present experiments, we test these hypotheses by probing readers' inferences about intended meanings of sentences and manipulating the experimental context to include sentences with specific types of errors (e.g., deletions, insertions, or exchanges). If comprehenders track the base-rate of errors but don't model the nature of the errors (context-invariant), they should make more inferences when they're exposed to errors than when the context contains only error-free sentences (Gibson et al., 2013), but the pattern of inferences should not differ by type of error exposure. However, if readers track more fine-grained error information beyond the base-rate (context-specific), their inferences should be sensitive to the type of error they experienced.

The goals of Experiment 1 were to a) replicate the effect of increasing the noise rate observed in Gibson et al. (2013) using a more direct measure (retyping and editing rather than comprehension questions), and b) test whether readers are sensitive to the nature of noise in the exposure. The goal of Experiment 2 was to run a pre-registered replication of Experiment 1 using a large sample size determined by a simulation-based power analysis of Experiment 1.