Contra assertions, feedback improves word recognition

James S. Magnuson^{1,2,3}, Samantha Grubb¹, Anne Marie Crinnion¹, Sahil Luthra¹, and Phoebe Gaston¹

¹University of Connecticut. Storrs, CT, USA

²BCBL. Basque Center on Cognition Brain and Language, Donostia-San Sebastián, Spain

³Ikerbasque. Basque Foundation for Science, Bilbao, Spain

Author Note

This research was supported by NSF BCS-PAC 1754284, NSF BCS-PAC 2043903, and NSF NRT 1747486 (PI: JSM). SL was supported by an NSF Graduate Research Fellowship. AMC and PG were supported by NIH T32 DC017703 (E. Myers and I-M. Eigsti, PIs). All stimuli, data, and analysis scripts are available at [this OSF repository](#). We thank Arty Samuel, Thomas Hannagan, and Rachel Theodore for comments that improved this work. Correspondence concerning this article should be addressed to James Magnuson, Department of Psychological Sciences, University of Connecticut, Storrs, CT, 06069-1020, USA. E-mail: james.magnuson@uconn.edu

Abstract

Norris and Cutler (in press) revisit their arguments that (lexical-to-sublexical) feedback cannot improve word recognition performance, based on the assumption that feedback must boost signal and noise equally. They also argue that demonstrations that feedback improves performance (Magnuson, Mirman, Luthra, Strauss, & Harris, 2018) in the TRACE model of spoken word recognition (McClelland & Elman, 1986) were artifacts of converting activations to response probabilities. We first evaluate their claim that feedback in an interactive activation model must boost noise and signal equally. This is not true in a fully interactive activation model such as TRACE, where the feedback signal does not simply mirror the feedforward signal; it is instead shaped by joint probabilities over lexical patterns, and the dynamics of lateral inhibition. Thus, even under high levels of noise, lexical feedback will selectively boost signal more than noise. We demonstrate that feedback promotes faster word recognition and preserves accuracy under noise whether one uses raw activations or response probabilities. We then document that lexical feedback selectively boosts signal (i.e., lexically-coherent series of phonemes) more than noise by tracking sublexical (phoneme) activations under noise with and without feedback. Thus, feedback in a model like TRACE does improve word recognition, exactly by selective reinforcement of lexically-coherent signal. We conclude that whether lexical feedback is integral to human speech processing is an empirical question, and briefly review a growing body of work at behavioral and neural levels that is consistent with feedback and inconsistent with autonomous (non-feedback) architectures.

Contra assertions, feedback improves word recognition

Introduction

Norris and Cutler (in press) make several excellent points in their homage to Jacques Mehler. These include the importance of a teleological approach (familiar to cognitive scientists as the *computational* level of Marr's taxonomy of information processing theories; Marr, 1982), the distinction between *using* ideal observer models and *assuming* optimality in biological systems (the ideal observer establishes a baseline against which to evaluate the system), and equivalence between Bayesian and interactive activation models under particular conditions (McClelland, 2013).

They then turn to the feedback debate in spoken word recognition, which has resisted resolution for decades. Recently, our team laid out a comprehensive case, via argument and demonstrations with model simulations, for how feedback can promote faster processing and make recognition robust in noise (Magnuson et al., 2018). In response, Norris, McQueen, and Cutler (2018) claimed that demonstrating a role for feedback in TRACE would be inconsequential (because it might not help in other models) and that arguments on the basis of parsimony were problematic without formal model comparisons. Feedback, they argued, should only be added to models of spoken word recognition if it is necessary to do so to account for human performance, and they emphasized that there was not conclusive evidence for feedback in behavioral studies with human subjects.

Since then, our team has provided what we believe to be conclusive behavioral evidence for feedback (Luthra, Peraza-Santiago, et al., 2021). We revisited the *lexically-mediated compensation for coarticulation* paradigm (Elman & McClelland, 1988) that parties on both sides of the debate agree provides a crucial test for feedback (Norris, McQueen, & Cutler, 2016; Pitt & McQueen, 1998), but has generated variable results (e.g., Magnuson, McMurray, Tanenhaus, & Aslin, 2003; McQueen, Jesse, & Norris, 2009; Samuel & Pitt, 2003). In a preregistered study, we first selected items only if they were able to drive the two component effects required for lexically-mediated compensation – lexical phoneme restoration (Ganong, 1980) and compensation for coarticulation (Mann & Repp, 1980, 1981). When we combined those items in the full paradigm, we found robust effects attributable to feedback in a well-powered sample ($n = 40$) and a direct replication ($n = 40$).

Norris and Cutler (in press) also introduce new assertions about feedback in TRACE, and in particular, claim that the simulations of Magnuson et al. (2018) were the result of a flawed "workaround" in TRACE. Here, we demonstrate via computational simulations that these assertions are incorrect. In an effort to avoid further misunderstandings, we provide detailed results from TRACE simulations that demonstrate how feedback promotes faster, more robust word recognition by

selectively reinforcing signal more than noise.

The new assertions

Here is the central claim made by Norris and Cutler (in press) regarding feedback:

In the case of speech perception, the most frequently encountered example is the claim that recognition benefits from interaction between different levels of processing, including feedback of activation from lexical to pre-lexical processing (“activation feedback”). Such interaction in fact serves no useful function at all and cannot improve recognition outcomes. (Norris & Cutler, in press, p. 3)

They then go on to make their case with a discussion of interactive activation and the TRACE model in particular. However, some crucial assertions are based upon flawed premises. Consider this passage.

Consider how feedback from lexical to phonemic processing is embodied in the TRACE model (McClelland & Elman, 1986); when a word is activated, that word sends activation back to the earlier stage (phoneme recognition) and activates the word’s constituent phonemes. In turn, these phonemes send yet more activation on to the lexical level. But what does this achieve? Activation from the (already best matching) lexical items has flowed back down to the phonemic representations and further increased their activation, and this can then further activate the same lexical items. Since these were already the best matching candidates, recognition does not improve; all that happens is reinforcement of the status quo. (Norris & Cutler, in press, p. 3)

This assertion follows from a misunderstanding of the dynamics in interactive activation models. Multiple features of interactive activation dynamics preclude the simple mirroring of feedforward and feedback flow. First, asymmetries are introduced by the nonlinear activation function and decay that is proportional to activation. Second, lateral inhibition can radically reorganize activation patterns in interacting layers. Third, and crucially, feedback can also change activation patterns on the inferior layer.

Figure 1 provides a schematic of a very simple interactive activation network with just seven phonemes and three words (*CAT*, *TELL*, *LOAD*). In this simple network, with no ability to encode temporal order, the input /kæt/ would strongly activate *CAT* and also weakly activate *TELL* (since /t/ connects to both words). Now consider what happens when noise is added to the initial input.

Figure 1. Interactive activation example. Seven input phonemes feed forward to three words, which send feedback to their constituent phonemes. Edges without arrows indicate lateral inhibition links within layers. Figure from Magnuson (2021).

Feedback will not simply reinforce the original feedforward pattern (of signal plus noise). For example, given the input */kæt/ + noise*, noise will likely provide more bottom-up input to *TELL* than the signal alone did. However, the noise will not be reinforced by feedback to the same extent as coherent patterns that map fully onto words (i.e., until the noise is very strong, it is unlikely that the constituent phonemes of a word dissimilar from the signal pattern will be as strongly activated *as a set* as the phonemes corresponding to the signal). Even slightly greater activation of *CAT* relative to *TELL* from the coherent input pattern to which noise has been added will allow *CAT* to inhibit *TELL* significantly (depending on exact parameters). Over iterations, this will lead *CAT* to send increasingly more feedback to its constituent phonemes relative to the feedback *TELL* sends to its constituent phonemes; *coherent patterns in the input will be reinforced to the detriment of nodes activated primarily by noise*, rather than simple "reinforcement of the status quo." In other words, the joint probability of {k, æ, t} embodied in the lexicon drives selective reinforcement of coherent lexical (and sublexical, phonotactically probable) patterns in the input. This will also drive faster and stronger activation of words consistent with the input, as we shall see in our TRACE simulations below.

The next passage reveals additional misunderstandings.

It is sometimes argued that activation feedback is required to ensure that recognition processes are more robust should speech be degraded. So, can recognition of a degraded signal be improved by the kind of activation feedback incorporated in TRACE? The answer

is again no. The above argument applies regardless of whether the input is degraded. The degraded-speech case still regularly causes confusion (e.g., Magnuson et al., 2018), and may have its origin in the original version of TRACE, which ruled out noise both in the input and during processing. (This was never intended as a theoretical claim about how speech recognition works; it just helped to simplify the model for computational purposes. The effect of noise was simulated by adding a constant amount of noise to a decision process – the Luce choice rule – operating on the output of the network.) In such a system, where feedback and noise are operating at different levels, feedback (as part of processing) can alter the relative activation of a word and its competitors but at the same time have no effect at all on any noise (which operates separately on final outputs). In other words, because of a workaround in the model, simulations using TRACE can give the impression that feedback can improve performance. (Norris & Cutler, in press, p. 3)

A key problem here seems to be about the use of the word *noise* to mean either uncertainty or actual addition of noise to a signal. The point of confusion appears to be the relation of the Luce (1959) choice rule with noise. The variant of the Luce choice rule employed in foundational work using interactive activation models (McClelland & Elman, 1986; McClelland & Rumelhart, 1981) may be familiar to readers as the *softmax* or *normalized exponential* function, with a weighting factor, k :

$$P_i = \frac{e^{ka_i}}{\sum_{j=1}^n e^{ka_j}} \quad (1)$$

Here, P_i is the response probability for choice i (i.e., word node i). The numerator calculates the *response strength* of node i : a_i is the activation of node i , and k is a constant that changes the degree of squashing (the larger k is, the more larger values will be amplified and the more smaller values will be damped). In the denominator, we calculate the summed response strengths for all n possible choices (including i). This might thus be considered an index of confidence, but it is not directly related to noise added to the signal. Consistent with Luce’s extensive work showing that such transformations provide better fits to human choice performance, TRACE activations transformed to response probabilities provide better fits to eye tracking data (e.g., Allopenna, Magnuson, & Tanenhaus, 1998).

The confusion may stem from the fact that McClelland (1991) explains that the original interactive activation models used the Luce choice rule to approximate (letter or phonemic) choice under uncertainty rather than actually adding noise to inputs or model-internal elements to simulate uncertainty. As Massaro (1989) pointed out, response probabilities calculated in this way fail to capture classical context effects in speech perception (e.g., continua inserted within frames with

different phonotactic probabilities, such as /l/ to /r/ within /s_i/ vs. /ʃ_i/). That is, the model fails to demonstrate what McClelland (2013) terms *logistic additivity* (essentially, independent effects of bottom-up and contextual information in frames such as those just described). As McClelland (1991) demonstrates, adding noise to raw activations (without the response probability conversion) and then averaging over many simulations *does* correctly capture classical context effects (see McClelland, 2013, for additional discussion and analysis).¹

Second, the previous quote suggests that Norris and Cutler (in press) believe that Magnuson et al. (2018) used the Luce choice rule to simulate the addition of noise. However, Magnuson et al. (2018) specify (p. 3) that we added Gaussian noise independently to each input element (following the procedure of McClelland, 1991). We subsequently converted activations to response probabilities in accordance with the procedures used by Frauenfelder and Peters (1998) to simulate lexical decisions (pp. 118-119), which were based on procedures used by McClelland and Rumelhart (1981) and McClelland and Elman (1986) for letter and phoneme recognition. Since our study was motivated by the investigation of feedback by Frauenfelder and Peters (1998) (cited by Norris, McQueen, & Cutler, 2000, without criticism of the use of response probabilities), and, as mentioned above, response probabilities generally provide better approximations of human performance, it was reasonable to follow their methods.

Third, consulting Equation 1, we can see that the impact of this simple transformation will be limited; it normalizes exponentially-transformed activations. It cannot, for example, change the rank order of choices. The claim that a "workaround in the model...can give the impression that feedback can improve performance" is very easy to test. A free, cross-platform version of TRACE (*jTRACE*; Strauss, Harris, & Magnuson, 2007) is freely available for download ². To test this claim, one could run four simulations with the default input /ʌbrʌpt/ and feedback set to 0.0 or the standard value of 0.03 and with input noise set to 0.0 or to have a moderate standard deviation (e.g., 0.75). Readers interested in conducting this test themselves with *jTRACE* (or the forthcoming *jsTRACE*; Crinnion, Grubb, Curtice, & Magnuson, in preparation) will find clear changes in phoneme and word activations with and without feedback *without applying the transformation to response probabilities*.

To provide a more robust test, we present simulations in the following section conducted under a range of input noise levels with feedback on and off. As we shall see, results with raw activations closely

¹ In practical use of TRACE, raw activations without noise added or simple response probabilities suffice for providing good fits to a variety of empirical phenomena; one might reasonably argue that input or internal noise (i.e., the stochastic version of TRACE McClelland, 1991) should be used, though the full scope of limitations of using the non-stochastic version has not been investigated, to our knowledge).

² <https://magnuson.psy.uconn.edu/jtrace/>

resemble those with response probabilities. We extend our previous results by tracking sublexical (i.e., phoneme) activations to demonstrate how feedback selectively reinforces lexically-coherent patterns.

Simulations

We use the same approach as Magnuson et al. (2018) to compare word recognition accuracy and recognition time in TRACE under increasing levels of noise with and without feedback – with the important alteration that we use raw activations rather than response probabilities.

Procedure

We conducted simulations using jsTRACE, a recent reimplementation of TRACE in JavaScript (Crinnion et al., in preparation). jsTRACE performs identically to jTRACE, but updates aspects of the Java implementation that are approaching obsolescence and provides the ability to do batch scripting via external JavaScript code. We used the default *slex* TRACE lexicon, consisting of 212 words (as well as the "silence" word used to represent a state of no input; the silence word was not included in analyses). We used three levels of feedback (0.00, 0.015, and 0.03, the last being the default level with small lexicons in TRACE). We combined each level of feedback with seven levels of Gaussian noise (with mean of zero and standard deviation ranging from 0.0 to 1.5 in steps of 0.25). A value sampled from the distribution was added independently to each cell of the input matrix prior to the simulation. Any negative input values were replaced with zero. To ensure that results under noise were robust, we conducted 10 simulations of every word in the lexicon at all levels of noise greater than zero. We allowed simulations to run for 100 time steps (cycles) in TRACE.

Decision policy. To assess how quickly a model "identifies" a word, we need a decision policy. We could choose the 'winner' to be the word that achieves maximum activation. However, how do we determine the *time* of recognition in that case? We cannot simply take the time of peak activation, as a target's activation may continue increasing throughout an entire simulation (thus resulting in longer recognition times for words that are more strongly activated, relative to a word that might have an earlier but lower peak). A further complication is that if we were to present the model with nonword (or novel word) inputs, simply taking the word with maximum activation as the winner would lead to erroneous 'recognition'. We used a simple threshold-based policy, where a correct identification is defined as the target reaching or exceeding that threshold and no other item reaching it. Recognition time is the cycle where the target's activation first reaches or exceeds the threshold. We first identified the activation threshold that would maximize accuracy for zero feedback without noise; this was 0.4. We then applied that threshold to every simulation (that is, at every level of feedback and noise). Note

that any potential bias in this policy favors simulations without feedback, since the threshold optimizes accuracy with zero noise and zero feedback.

Crucially, all analyses were applied to raw activations. We did not transform activations to response probabilities.

Figure 2. Replication of Magnuson et al. (2018) using activations instead of response probabilities. Each point represents the outcome of simulating every word in the 212-word *slx* lexicon, with 10 simulations conducted with each word at each noise level greater than zero. The threshold was set to 0.4, which maximized accuracy without feedback and with noise set to zero.

Results

In Figure 2, we see a clear replication of the results of Magnuson et al. (2018) based on raw activations. Recognition time is faster with feedback than without at every level of noise – with the exception of $sd = 1.25$, where only 10 correct trials (out of 2120 trials) occurred with feedback set to 0.0 (compared to 126 correct trials with feedback at 0.015 and 290 correct trials with feedback at 0.03), making this datapoint virtually uninterpretable. Feedback yields higher accuracy once noise is added.

In Figure 3, we plot recognition times for words that were correctly recognized with feedback at the default value of 0.03 and without feedback (feedback set to 0.0) at each noise level. Again, we repeated the simulation of each word 10 times at each level of noise greater than zero. A clear, consistent advantage is observed for the majority of words at each level of noise.

Contra assertions by Norris and Cutler (in press), feedback promotes more robust word recognition performance, and therefore the results reported by Magnuson et al. (2018) were not an artifact of using

Figure 3. Comparing activation-based recognition time with feedback (set to 0.03) and without feedback (set to 0.0). Only words that were correctly recognized both with and without feedback are included. Without noise (top left), we conducted one simulation per word with and without feedback (without noise, the outcome would be identical if we repeated simulations); 196 of 212 words are included because those were correctly recognized both with and without feedback. At each noise level greater than 0 (all other panels), there were 2120 simulations (10 repetitions of each word with Gaussian noise added to the input). Items classified as "faster" were recognized more quickly (reached the threshold more quickly) with feedback than without, and are below the identity line; "equal" items are on the identity line; "slower" items reached the threshold later with feedback than without, and are above the identity line.

response probabilities rather than raw activations. TRACE recognized words more quickly with feedback than without at every level of noise (including no noise), and feedback promoted higher accuracy as noise was added to simulations. How would this be possible if feedback simply reflected signal and noise equally, as Norris and Cutler (in press) claim? We can explore this by examining how phoneme and word activations change as noise is added with and without feedback.

First, consider phoneme activations (Figure 4). In the figure, we plot the mean activation of each input phoneme (solid red lines) averaged over all items (all words in *slex*) along with the mean activation of the next-most-activated phoneme at that position (dashed blue lines). Although the maximum word length is 9, the number of words contributing to means for later phonemes decreases, since most words are shorter. There is a clear pattern: with feedback (right panels), there is greater separation between Target phonemes and Next phonemes than without (left panels). We can see this is a result both of greater Target phoneme activation with feedback vs. without and modestly lower activation of Next phonemes with feedback vs. without (though this is subtle, and mainly observable as sharper peaks for Next phonemes with feedback, indicative of greater lateral inhibition from Target phonemes). If signal and noise were equally reinforced by feedback, we would expect to see similar amplification of both Target and Next phonemes with feedback. Consistent with our discussion of Figure 1, feedback selectively boosts those phonemes that constitute lexically-coherent patterns – that is, series of phonemes that constitute words.

Next, let’s consider what happens at the lexical level. In Figure 5, we present the mean activations of correctly-recognized target items with and without feedback as noise is added. For each item, we also tracked activations of *cohorts* (words sharing the first two onset phonemes) and *rhymes* (not poetic rhymes or linguistic rimes, but words mismatching by a single phoneme at onset position by deletion [*cat-at*], addition [*cat-scat*], or substitution [*cat-bat*]; cf. Allopenna et al., 1998).

All other words (other than targets, cohorts, or rhymes) were averaged to provide the *unrelated* baseline. We also tracked the maximally-activated non-target activation at every time step, labelled *OtherMax*. We did this because the *mean* cohort activation, for example, would be less than the *maximum* cohort activation; tracking the maximally-activated non-target (whether it is a cohort, rhyme, or other item) gives a better sense of the relative activation of the target.³

³ Although the result for *OtherMax* is a smooth trajectory (since we are averaging over simulations with potentially as many as 212 target words and 10 repetitions, depending on how many words were correctly recognized at a given combination of feedback and noise; see Figure 3 for numbers of items included) note that *OtherMax* does not correspond to a single item for any simulation, because it takes the maximal non-target activation at each time step, and which non-target word achieved maximal activation could change at each time step. Note also that items in *OtherMax* also contributed to other non-target categories (e.g., most *OtherMax* items were the most strongly activated cohort in a given simulation).

Figure 4. Phoneme activations over time for words at two levels of feedback and five levels of noise averaged over all 212 words in the *slax* lexicon. In each plot, the activations of the phoneme corresponding to the input (solid red lines) and the next-most activated phoneme (dashed blue lines) are plotted. Later phonemes are based on fewer words (those up to nine phonemes in length).

Figure 5. Lexical activations by category for two levels of feedback and five levels of noise, restricted to words that were correctly recognized at each level at a threshold of 0.4 (compare to accuracy in Figure 2). *OtherMax* is the mean activation of the most highly-activated non-target item at each time step. Standard error ribbons are plotted, but are too small to discern in most panels (but are apparent in a few cases, such as rhymes in the bottom-right panel). Rhyme lines are absent in some Feedback = 0 panels because the very small numbers of correctly recognized items (e.g., only 10 at Noise = 1.25) did not have rhymes. The panel for Feedback = 0, Noise = 1.5 is blank because no words were recognized at this combination.

For correctly-recognized word inputs (Figure 5), feedback is associated with somewhat higher target activations as noise increases, as well as more robust activation of phonological competitors (consistent with stronger activation of constituent phonemes). The differences are most apparent in higher mean target peaks with feedback than without when noise is equal to or greater than 1.

Discussion

The simulations we presented demonstrate clear benefits of feedback in the interactive framework. Feedback promotes faster recognition time and makes the model robust against noise, as demonstrated by preservation of accuracy with feedback as compared to without. Contra assertions made by Norris and Cutler (in press), this is not an artifact of using response probabilities rather than raw activations. Magnuson et al. (2018) used the standard approach of using response probabilities, but here we used raw activations.

One objection we anticipate is that we have shown a benefit that is specific only to the TRACE model. Indeed, Norris et al. (2018) make precisely this claim, and say the only reason feedback helps in TRACE is "because its initial behavior is suboptimal". Although we disagree about "suboptimality" here, we will argue, in the spirit of ideal observer models, that the issue of optimality is secondary. Suboptimal behavior is not necessarily *incorrect* behavior; the interesting question is to what degree a model corresponds with humans at behavioral, algorithmic, and neural levels (Magnuson et al., 2018), and optimality is an interesting baseline to consider.

Norris et al. (2018) argue that "...the best that any speech recognition system can do is compute the match between input features and lexical representations and select the best-matching word (more specifically, pick the word with maximum likelihood) ... Shortlist B...by virtue of implementing Bayesian inference, performs optimally; its use of Bayes' rule guarantees that the best-matching word must be recognized (p. 1)". Some details here may help us better understand and possibly refine the debate between autonomous (without feedback) and interactive (with feedback) accounts.

First, one question is *how* a system can compute the match between inputs and words. Shortlist B does this by consulting lookup tables. Empirically elicited diphone confusions are used to estimate probability distributions over phonemes at each position, which are then used to identify the word with highest likelihood. As a computational-level ideal observer theory, it is concerned with *why/should* questions rather than *how* (i.e., the mechanistic details of the word recognition process). It is concerned with the (optimal) mapping from inputs to outputs, that is, from phonemes to maximally likely words. As Norris and Cutler (in press) put it, an ideal observer model provides an explicit prediction of how an organism *should* behave, allowing one to mine actual behavior for deviations from

the model’s definition of optimality and, ideally, insights into the actual mechanisms at work.

TRACE, in contrast, is an algorithmic model (in the larger computational framework of *interactive activation*; McClelland, Mirman, Bolger, & Khaitan, 2014), and grapples deeply with the mechanistic details of *how* spoken word recognition might operate. It implicitly, rather than explicitly, computes matches between inputs and words. Rather than consulting a lookup table, complex, nonlinear interactions among levels of representation (features, phonemes, words) drive the model to activate words proportionally to likelihoods that are embodied in the contents of the lexicon. As has been discussed in detail elsewhere (e.g., Magnuson et al., 2018; McClelland, 2013), the dynamics of interactive activation provide a generative model; feedback reinforces sublexical activation patterns that are the most likely causes of lexical activations (this is apparent in the figures we have presented here, especially Figure 4).

Second, since interactive activation provides TRACE with a method to compute the match between inputs and words, we can ask two questions regarding optimality: (1) can a model like TRACE generate Bayes optimal behavior, and (2) is human word recognition Bayes optimal? The answer to the first question is *yes*. McClelland (2013) explains how the *multinomial interactive activation* (MIA) model is Bayes equivalent, although this requires some neurally-unrealistic constraints (e.g., making activation passing serial to emulate Gibbs sampling). However, McClelland (1991) and Movellan and McClelland (2001) demonstrate multiple ways that interactive activation models can produce logistically additive behavior (and thus simulate classical context effects on perceptual judgments); the simplest is just adding noise to inputs, as we have done in our simulations here.

Logistic additivity is more important, in our view, than strict Bayes optimality, as this is a key requirement for simulating aspects of human performance, whereas it is unclear that human performance is Bayes optimal, or even what the proper definition of *optimality* is. For example, humans often probability match rather than probability maximize (that is, apparently stochastically sampling among choices according to their likelihoods, rather than simply choosing the most likely choice). As Movellan and McClelland (2001) discuss, such deviations may reflect non-obvious maximization functions, such as the tension between *exploration* and *exploitation*. If an organism always maximizes, but probabilities in the environment are not stationary, the organism may ultimately make suboptimal choices. Favoring exploration (probability matching) over exploitation (maximizing) may ultimately be the more optimal strategy under non-stationary (i.e., realistic) priors. Furthermore, aspects of human perception and cognition may approximate optimality despite relying on heuristic or other non-Bayesian mechanisms (Gigerenzer & Goldstein, 1991; Hecce Castañón et al., 2019; Rahnev & Denison, 2018; Zhou & Yu, 2021).

Biological systems may evolve mechanisms that deviate from normative Bayesian optimality in order to maximize efficiency given limited neural, sensory, and cognitive resources (Griffiths, Lieder, & Goodman, 2015; Stengård & van den Berg, 2019; Summerfield & Li, 2018), or that "satisfice" to achieve "good enough" performance under constraints of *bounded rationality* rather than optimality (Gigerenzer & Goldstein, 1991; Simon, 1956). Indeed, there are many examples where human language processing appears to deviate from optimality, such as rhyme effects in spoken word recognition even when word onsets are perfectly clear (Alloppenna et al., 1998), activation of additional potential agents or patients even when a role is already filled (Kukona, Fang, Aicher, Chen, & Magnuson, 2011), local coherence effects in sentence processing (Tabor, Galantucci, & Richardson, 2004), and simultaneous inference of incompatible sentences parses (Ferreira, Bailey, & Ferraro, 2002). The question of whether a model is Bayes optimal, Bayes approximate, or systematically suboptimal is fundamentally interesting, but is secondary to the nature of the biological capacity that is being modeled. Ideal observer models are useful whether or not the capacity being modeled is optimal, but are wrong when the capacity is not optimal. The crucial question is whether the model is qualitatively like the capacity being modeled; does it, for example, achieve, approximate, or deviate from optimality as the biological system does?

Thus, a vital concern is a model's ability to simulate the capacity being modeled in detail, and to inform theories at Marr's (1982) computational, algorithmic, and implementational levels. Not all models or theories connect to all three levels, but we agree with Marr that comprehensive understanding requires accounts at all three levels. Leaving aside questions of optimality, TRACE has contributed significantly at all three levels. Magnuson and Crinnion (in press) tabulated 30 specific empirical phenomena of speech perception and spoken word recognition that have been modeled with TRACE (ranging from categorical perception to phoneme restoration and the time course of phonological competition), compared to only 7 with Shortlist B. In addition, TRACE is proving surprisingly capable of simulating neural level data, including crucial patterns of neural activity hypothesized to be linked to predictive coding. Luthra, Li, You, Brodbeck, and Magnuson (2021) demonstrate that TRACE captures a key pattern observed by Gagnepain, Henson, and Davis (2012) in fMRI: signal reduction when inputs conform to expectations. Brodbeck, Luthra, Gaston, and Magnuson (2021) demonstrate that TRACE captures a complex cross-over interaction observed in representational similarity analyses (Kriegeskorte & Bandettini, 2008) of fMRI (Blank & Davis, 2016) and MEG (Sohoglu & Davis, 2020) data (and demonstrate how this depends crucially on feedback), and also document interesting similarities between MEG studies of speech processing and the sensitivity of different levels of representation in TRACE to information theoretic measures such as surprisal and entropy (and how those relate to specific aspects of TRACE's architecture, such as lateral inhibition and feedback).

Conclusions

The purposes of models like TRACE and Shortlist B are qualitatively different in important ways. As Norris and Cutler (in press) review, Shortlist B is an *ideal observer* model. It models optimal inference, based upon human diphone confusion probabilities (its inputs) and Bayes' rule. By comparing human performance to it, one may observe how human performance differs from its definition of optimality. Those differences could be used to inform a more realistic theory of human word recognition (although so far as we are aware there no such theory has emerged yet). It is situated at Marr's (1982) *computational* level of information processing theories. Its relation to algorithmic and implementational levels is unclear, though Norris et al. (2016) suggest that implementational details should have little bearing on computational level theories (p. 7).

TRACE is situated primarily at Marr's (1982) *algorithmic* level of information processing theories, but it connects to computational and implementational questions as well. It extends the interactive activation computational theory to the domain of speech and spoken word recognition. Its focus is explicitly not just *what* the mapping is from inputs to word forms, but *how* that mapping can be achieved within a framework that is explicitly motivated by basic elements of neural systems.⁴ As such, it is also poised to bridge perceptual, cognitive, and neural studies of language processing. It is obviously useful for researchers interested in making those connections (as we discussed above, TRACE is being used productively to investigate neural data: Brodbeck et al., 2021; Luthra, Li, et al., 2021).

While we think it is important to correct the record when it comes to false assertions in the literature (our motivation in engaging in this paper), we think debates about optimality and whether feedback is logically necessary have run their course, and extending them offers little further utility. It will be more productive to engage in detailed assessments of various models' abilities to explain computational, behavioral, developmental, and neural aspects of biological systems, to develop new models, and to compare models to each other – not with the goal of finding the "winner", but with the goal of furthering understanding of the cognitive, perceptual, and neurobiological bases of human language processing by discovering how various models deviate from experimental evidence, progressively revealing more constraints on the underlying human system.

⁴ Norris and Cutler (in press) argue that theorists must be concerned with *why* theories are structured as they are, not just *how* a system works. The originators of interactive activation and related parallel distributed processing theories and models certainly did this, with an entire volume focused on *why* (Rumelhart, McClelland, & the PDP Research Group, 1986) and another on *how* (McClelland, Rumelhart, & the PDP Research Group, 1986).

References

- Alloppenna, P., Magnuson, J., & Tanenhaus, M. (1998). Tracking the time course of spoken word recognition using eye movements: Evidence for continuous mapping models. *Journal of Memory and Language*, 38, 419–439.
- Blank, H., & Davis, M. H. (2016). Prediction errors but not sharpened signals simulate multivoxel fmri patterns during speech perception. *PLoS Biology*, 14(11), 1–33.
- Brodbeck, C., Luthra, S., Gaston, P., & Magnuson, J. S. (2021). Discovering computational principles in models and brains. *Proceedings of the Cognitive Science Society*.
- Crinnion, A. M., Grubb, S., Curtice, A., & Magnuson, J. S. (in preparation). jsTRACE: A javascript reimplimentation of the TRACE model of speech perception and spoken word recognition.
- Elman, J. L., & McClelland, J. L. (1988). Cognitive penetration of the mechanisms of perception: Compensation for coarticulation of lexically restored phonemes. *Journal of Memory and Language*, 27(2), 143–165.
- Ferreira, F., Bailey, K. G., & Ferraro, V. (2002). Good-enough representations in language comprehension. *Current Directions in Psychological Science*, 11(1), 11–15. doi: 10.1111/1467-8721.00158
- Frauenfelder, U. H., & Peters, G. (1998). Simulating the time course of spoken word recognition: An analysis of lexical competition in TRACE. In *Localist connectionist approaches to human cognition*. (pp. 101–146). Mahwah, NJ, US: Lawrence Erlbaum Associates Publishers.
- Gagnepain, P., Henson, R. N., & Davis, M. H. (2012). Temporal predictive codes for spoken words in auditory cortex. *Current Biology*, 22(7), 615–621.
- Ganong, W. F. (1980). Phonetic categorization in auditory word perception. *Journal of Experimental Psychology: Human Perception and Performance*, 6(1), 110–125. doi: <https://doi.org/10.1037/0096-1523.6.1.110>
- Gigerenzer, G., & Goldstein, D. G. (1991). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4), 650–669.
- Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, 7(2), 217–229.
- Herce Castañón, S., Moran, R., Ding, J., Egner, T., Bang, D., & Summerfield, C. (2019). Human noise blindness drives suboptimal cognitive inference. *Nature Communications*, 10(1719). doi: 10.1038/s41467-019-09330-7
- Kriegeskorte, M. M., N., & Bandettini, P. A. (2008). Representational similarity analysis—connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2, 1–28.
- Kukona, A., Fang, S., Aicher, K. A., Chen, H., & Magnuson, J. S. (2011). The time course of

- anticipatory constraint integration. *Cognition*, 119, 23-42.
- Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. New York, USA: Wiley.
- Luthra, S., Li, M. Y. C., You, H., Brodbeck, C., & Magnuson, J. S. (2021). Does signal reduction imply predictive coding in models of spoken word recognition? *Psychonomic Bulletin and Review*. doi: 10.3758/s13423-021-01924-x
- Luthra, S., Peraza-Santiago, G., Beeson, K., Saltzman, D., Crinnion, A. M., & Magnuson, J. S. (2021). Robust lexically-mediated compensation for coarticulation: Christmash time is here again. *Cognitive Science*, 45(4), e12962.
- Magnuson, J. S. (2021). Interactive activation example network. Retrieved from <https://doi.org/10.6084/m9.figshare.14829078.v1>
- Magnuson, J. S., & Crinnion, A. M. (in press). Spoken word recognition. In J. Trueswell, A. Papafragou, & L. Gleitman (Eds.), *Oxford handbook of the mental lexicon*.
- Magnuson, J. S., McMurray, B., Tanenhaus, M. K., & Aslin, R. N. (2003). Lexical effects on compensation for coarticulation: The ghost of christmash past. *Cognitive Science*, 27, 285–298.
- Magnuson, J. S., Mirman, D., Luthra, S., Strauss, T., & Harris, H. (2018). Interaction in spoken word recognition models: Feedback helps. *Frontiers in Psychology*, 9(369). doi: 10.3389/fpsyg.2018.00369
- Mann, V., & Repp, B. (1980). Influence of vocalic context on perception of the [ʃ]–[s] distinction. *Perception Psychophysics*, 28, 213–228.
- Mann, V., & Repp, B. (1981). Influence of preceding fricative on stop consonant perception. *Journal of the Acoustical Society of America*, 69, 548–558.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. USA: Henry Holt and Co. , Inc.
- Massaro, D. W. (1989). Testing between the TRACE model and the fuzzy logical model of speech perception. *Cognitive Psychology*, 21(3), 398–421. doi: 10.1016/0010-0285(89)90014-5
- McClelland, J. L. (1991). Stochastic interactive processes and the effect of context on perception. *Cognitive Psychology*, 23, 1-44. doi: 10.1016/0010-0285(91)90002-6
- McClelland, J. L. (2013). Integrating probabilistic models of perception and interactive neural networks: a historical and tutorial review. *Frontiers in Psychology*, 4, 503. doi: 10.3389/fpsyg.2013.00503
- McClelland, J. L., & Elman, L., J. (1986). The TRACE model of speech perception. *Cognitive Psychology*, 18, 1-86. doi: 10.1016/0010-0285(86)90015-0
- McClelland, J. L., Mirman, D., Bolger, D. J., & Khaitan, P. (2014). Interactive activation and mutual constraint satisfaction in perception and cognition. *Cognitive Science*, 6, 1139-1189. doi: 10.1111/cogs.12146

- McClelland, J. L., Rumelhart, D., & the PDP Research Group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. volume i*. Cambridge, MA: MIT Press.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review*, 88(5), 375–407. (Place: US Publisher: American Psychological Association) doi: 10.1037/0033-295X.88.5.375
- McQueen, J. M., Jesse, A., & Norris, D. (2009). No lexical-prelexical feedback during speech perception or: Is it time to stop playing those christmas tapes? *Journal of Memory and Language*, 61(1), 1–18.
- Movellan, J. R., & McClelland, J. L. (2001). The morton-massaro law of information integration: Implications for models of perception. *Psychological Review*, 108(1), 113–148. doi: 10.1037/0033-295X.108.1.113
- Norris, D., & Cutler, A. (in press). More why, less how: What we need from models of cognition. *Cognition*.
- Norris, D., McQueen, J. M., & Cutler, A. (2000). Merging information in speech recognition: feedback is never necessary. *Behavioral and Brain Sciences*, 23(3), 299–370. doi: 10.1017/s0140525x00003241
- Norris, D., McQueen, J. M., & Cutler, A. (2016). Prediction, Bayesian inference and feedback in speech recognition. *Language, Cognition and Neuroscience*, 31(1), 4–18.
- Norris, D., McQueen, J. M., & Cutler, A. (2018). Commentary on “interaction in spoken word recognition models”. *Frontiers in Psychology*, 9, 1–3. doi: 10.3389/fpsyg.2018.01568
- Pitt, M. A., & McQueen, J. M. (1998). Is compensation for coarticulation mediated by the lexicon? *Journal of Memory and Language*, 39, 347–370.
- Rahnev, D., & Denison, R. (2018). Suboptimality in perceptual decision making. *Behavioral and Brain Sciences*, 41(e223), 1–66. doi: 10.1017/S0140525X18000936
- Rumelhart, D., McClelland, J. L., & the PDP Research Group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. volume ii*. Cambridge, MA: MIT Press. [Archive.
- Samuel, A. G., & Pitt, M. A. (2003). Lexical activation (and other factors) can mediate compensation for coarticulation. *Journal of Memory and Language*, 48(2), 416–434.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129–38.
- Sohoglu, E., & Davis, M. H. (2020). Rapid computations of spectrotemporal prediction error support perception of degraded speech. *ELife*, 9, e58077.
- Stengård, E., & van den Berg, R. (2019). Imperfect bayesian inference in visual perception. *PLOS Computational Biology*, 15(4), 1–27. doi: 10.1371/journal.pcbi.1006465
- Strauss, T. J., Harris, H. D., & Magnuson, J. S. (2007). jTRACE: A reimplementaion and extension

- of the TRACE model of speech perception and spoken word recognition. *Behavior Research Methods*, 39, 19-30.
- Summerfield, C., & Li, V. (2018). Perceptual suboptimality: Bug or feature? *Behavioral and Brain Sciences*, 41(e245), 39-40. doi: 10.1017/S0140525X18001437
- Tabor, W., Galantucci, B., & Richardson, D. (2004). Effects of merely local syntactic coherence on sentence processing. *Journal of Memory and Language*, 50(4), 355-370.
- Zhou, Y., & Yu, Y. (2021). Human visual search follows a suboptimal Bayesian strategy revealed by a spatiotemporal computational model and experiment. *Communications Biology*, 4(34). doi: 10.1038/s42003-020-01485-0