An information-theoretic account of availability effects in language production

Richard Futrell
Department of Language Science
University of California, Irvine
rfutrell@uci.edu

Abstract

I present a computational-level model of language production in terms of a combination of information theory and control theory in which words are chosen incrementally in order to maximize communicative value subject to an information-theoretic capacity constraint. The theory generally predicts a tradeoff between ease of production and communicative accuracy. I apply the theory to two cases of apparent availability effects in language production, in which words are selected on the basis of their accessibility to a speaker who has not yet perfectly planned the rest of the utterance. Using corpus data on English relative clause complementizer dropping from Levy & Jaeger (2007) and experimental data on Mandarin noun classifier choice from Zhan & Levy (2019), I show that the theory reproduces the observed phenomena, providing an alternative account to Uniform Information Density (UID) and a promising general model of language production which is tightly linked to emerging theories in computational neuroscience.

Keywords: language production, information theory, control theory, availability-based production, accessibility

Introduction

Language production appears to be a largely incremental process: speakers plan an utterance as they are producing it, simultaneously integrating multiple sources of information (Bock, 1982; Levelt, 1989; Bock & Levelt, 1994; Griffin, 2001; F. Ferreira & Swets, 2002). One apparent effect of this incrementality is **availability** effects in language production: the fact that speakers will often choose to produce words which are easily accessible or available to them earlier in an utterance, or to include such words when they are optional, or even to use a highly-available word in place of a more communicatively accurate but less available word (Koranda et al., 2021). The fact that available words tend to go earlier has been attributed to a greedy, 'easy-first' language production strategy (Bock & Irwin, 1980; Levelt, 1981; Bock, 1982; Bock & Warren, 1985; McDonald et al., 1993; V. S. Ferreira & Dell, 2000; V. S. Ferreira & Yoshita, 2003; Chang, 2009; Tanaka et al., 2011; F. Ferreira & Rehrig, 2019).

Here I present an account of availability effects within a computational-level model of language production based on a recently developing theory of the complexity of action selection from the fields of computational neuroscience and information theory. This theory, the **rate–distortion theory of control** (RDC), holds that actions are selected to maximize value subject to constraints on the use of information. The theory originates in the economics literature (C. A. Sims, 2003) where it operationalizes bounded rationality (Simon, 1955; Lewis et al., 2014; Lieder & Griffiths, 2019), and it has been developed and applied in the literature on physics, robotics, optimal control, computational neuroscience, reinforcement learning, cognitive psychology, and linguistics (Rubin et al., 2012; Tishby & Polani, 2011; Todorov, 2009; Van Dijk & Polani, 2013; Genewein et al., 2015; Ortega & Braun, 2013; Gershman, 2020; Gershman & Bhui, 2020; Lai & Gershman, 2021; Bhui et al., 2021; C. R. Sims, 2016, 2018; Zaslavsky et al., 2021; Arumugam et al., 2023). RDC uses the mathematical theory of lossy compression (Shannon, 1959; Cover & Thomas, 2006) to impose informational constraints on the perception–action loop. It has also been termed rational inattention and policy compression.

I develop a proof-of-concept model of language production within the RDC framework, based on an informational constraint identifiable as a channel capacity limit on cognitive control (Fan, 2014; Zénon et al., 2019). I show that this model provides an account of availability effects in language production, and I validate this account by examining experimental data from two previous sets of experiments: Levy & Jaeger (2007) on relative clause complementizers in English, and Zhan & Levy (2019) on noun classifier choice in Mandarin Chinese. In contrast with existing models of language production which are primarily situated at Marr's (Marr, 1982) algorithmic level of analysis or at more concrete levels, the RDC model is at the computational level: it directly describes the inputs, outputs, and goals of the language production system, without committing to an algorithmic implementation. The high level of abstraction means that it is possible to see how simple underlying computational constraints give rise to a variety of different behaviors.

Model

Setting

The goal of the RDC production model is to characterize a speaker's **production policy**, which is the probability distribution over **actions** a (outputs) at each time, given the speaker's current **state** s and **goal** g (inputs), notated $\pi_g(a \mid s)$. An example of goal and state is shown in Figure 1A. In the studies below, an 'action' refers to the production of an individual word; I leave open the possibility that a model of the same form could apply for units as small as phonemes or as large as phrases.

In the studies below, the state s consists of the speaker's memory for what she has produced so far. After production of an action a, the state s updates to a new state s' and production continues with the selection of a new action conditional on s'. For simplicity, I assume below that the state contains a lossless memory of all actions produced so far in an utterance.

The communicative goal g represents the speaker's intent for an utterance. Following probabilistic models of semantics and pragmatics (Kemp & Regier, 2012; Frank & Goodman,

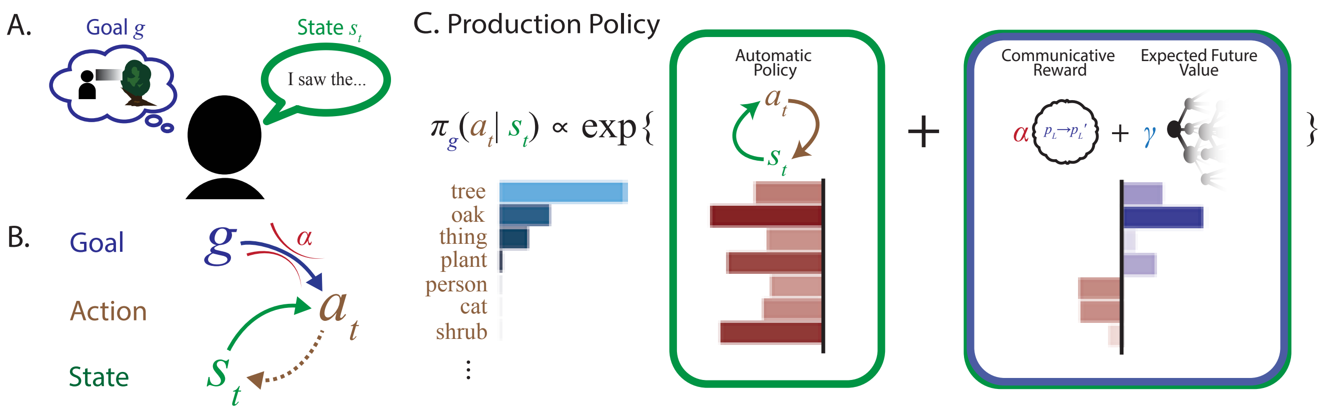

Figure 1: Schematic structure of the language production policy. During production of an utterance, at time t, an **action** (such as a word) a_t is selected probabilistically as a function of the **communicative goal** g and the **current state** s_t; after production of an action a_t, the state updates in response, and the cycle continues. The RDC model maximizes average communicative reward subject tox a channel capacity constraint on the use of control information in action selection, determined by the control gain parameter α. The resulting optimal controlled policy $\pi_g(x \mid s)$ combines (1) an **automatic policy**, representing what the system's behavior regardless of goal input, (2) the **communicative reward** with respect to the current goal g, defined as the amount of information provided to a listener model p_L, plus (3) the discounted expected value of subsequent actions.

2012; Goodman & Frank, 2016; Scontras et al., 2021; Zaslavsky et al., 2018; Mollica et al., 2021), I hold that the goal of producing a linguistic utterance is to communicate a world state to a listener. A communicative goal g thus specifies a target world state notated w_g.

Intuition

The key intuition that I use when applying RDC to language production is that it is costly for the speaker to integrate information about goals. Thus the policy is selected to maximize communicative reward subject to a constraint on how much information about the goal may be used at each timestep, as illustrated in Figure 1B. This constraint makes language production 'good enough': a speaker may produce a communicatively sub-optimal utterance if the optimal utterance would involve too much information-processing cost (F. Ferreira et al., 2002; V. S. Ferreira & Griffin, 2003; F. Ferreira & Patson, 2007; F. Ferreira & Lowder, 2016; Koranda et al., 2021; Goldberg & Ferreira, 2022). The model is domain-general in the sense that policies of the same form have been used to study general decision making and action planning (Tishby & Polani, 2011; Ortega & Braun, 2013; Van Dijk & Polani, 2013; Schach et al., 2018; Gershman, 2020; Gershman & Bhui, 2020; Bhui et al., 2021; Lai & Gershman, 2021).

The policy obtained from this constrained optimization problem has a dual structure in which two forces influence action selection. The first force favors selection of the most communicatively rewarding action according to the current goal (maximizing information conveyed about the speaker's target world state), or the action which leads to the most rewarding actions in the future. The second force favors selection of whatever action is most probable in context, without regard to the current goal. This second force captures automaticity in language production: frequent sequences of actions are produced 'automatically' and are hard to disrupt (Lounsbury, 1954; Shiffrin & Schneider, 1977; Kapatsinski, 2010).

Formal specification

The RDC language production model is a policy that maximizes a tradeoff of reward and informational complexity within a Markov Decision Process describing the production of an utterance. It is formally an instance of KL-control (Todorov, 2009). The policy $\pi_g(\cdot \mid s)$ for a given state s and goal g is selected to maximize a **value function** which is the average immediate reward for actions taken in that state, plus the expected future reward for following the same policy from the resulting state (Sutton & Barto, 2018):

$$\underbrace{V_g^\pi(s)}_{\text{Value}} = \mathop{\mathbb{E}}_{a\sim\pi_g(\cdot|s)}\left[\underbrace{\ell_g^\pi(a \mid s)}_{\text{Immediate value}} + \underbrace{\gamma V_g^\pi(s')}_{\text{Future value}}\right], \tag{1}$$

where s' is the state after production of action a, and the **future-discount parameter** $\gamma \in [0,1]$ determines how the value of future actions is weighted relative to current actions. The **immediate value** $\ell_g^\pi(a \mid s)$ of an action a given state s and goal g decomposes into communicative reward minus information-processing cost, called control cost:

$$\underbrace{\ell_g^\pi(a \mid s)}_{\text{Immediate value}} = \underbrace{\alpha R_g(a \mid s)}_{\text{Reward}} - \underbrace{\ln\frac{\pi_g(a \mid s)}{\pi_0(a \mid s)}}_{\text{Control cost}}. \tag{2}$$

Control cost represents a cost for using information about the goal g: it is the divergence between the **controlled policy** π_g, which uses information about the goal g, and a policy π_0 which uses no information about the goal g, called the **automatic policy**. Actions that deviate from the automatic policy are relatively costly to an extent that depends on the scalar **control gain** α which adjusts the balance between communicative reward and control cost. The control cost is the (pointwise) mutual information between actions and goals given states, so by penalizing control cost, we can limit the amount of information used about the goal to select each action.

I define the **communicative reward** $R_g(a \mid s)$ of an action a as the incremental information provided by action a about w_g for a listener L:

$$\underbrace{R_g(a \mid s)}_{\text{Reward}} = \frac{p_L(w_g \mid a, s)}{p_L(w_g \mid s)}, \tag{3}$$

where the distribution p_L is the **listener model**, a distribution on world states given utterances, representing the speaker's beliefs about how a listener will react to her utterances. Knowledge of language is localized within the listener model p_L which defines a stochastic mapping between utterances and world states. Unlike Rational Speech Acts (RSA) models of pragmatics (Frank & Goodman, 2012; Goodman & Lassiter, 2014; Scontras et al., 2021; Zaslavsky et al., 2021), the listener model p_L is not necessarily a Bayesian inversion of a speaker policy—it can be seen as an RSA literal listener.

The RDC policy is the distribution $\pi_g(a \mid s)$ on actions given states and goals that maximizes the average value as defined in Eq. 1. This policy has the form

$$\underbrace{\pi_g(a \mid s)}_{\text{Controlled policy}} \propto \exp\left\{ \underbrace{\ln \pi_0(a \mid s)}_{\text{Automatic policy}} + \underbrace{\alpha R_g(a \mid s)}_{\text{Reward}} + \underbrace{\gamma V_g^{\pi}(s')}_{\text{Future value}} \right\}, \tag{4}$$

as illustrated in Figure 1C. Thus, the choice to produce an action a is influenced by three factors, as shown in Figure 1: (1) the automatic policy π_0, (2) the communicative reward R_g, and (3) planning, in the form of the future value.

The automatic policy π_0 represents the actions that speaker is likely to take regardless of goal. Because it reflects the predictability of an action (word) given a state (context), it is a language model in the natural language processing sense (Andreas, 2022). The automatic policy can be derived from the controlled policy by marginalizing out the communicative goals:

$$\pi_0(a \mid s) = \sum_g p(g \mid s)\pi_g(a \mid s). \tag{5}$$

Eq. 5 depends on a **need distribution** on communicative goals $p(g)$ (Kemp & Regier, 2012; Kemp et al., 2018; Zaslavsky et al., 2020). To find a concrete controlled policy $\pi_g(a \mid s)$, one must start from a randomly-initialized controlled policy and repeatedly evaluate Eqs. 1, 4, and 5 until reaching a fixed point, in a form of value iteration (Sutton & Barto, 2018). This is because the value function in Eq. 1, the controlled policy in Eq. 4, and the automatic policy in Eq. 5 depend on each other recursively.

The RDC policy can be interpreted as maximizing reward subject to a constraint on usage of information about the goal (Rubin et al., 2012; Van Dijk & Polani, 2013), corresponding to a channel capacity constraint on cognitive control (Fan, 2014; Zénon et al., 2019). In this interpretation, cognitive control transmits signals at a limited rate in terms of bits per timestep. This bit-rate is related to the control gain α: higher α indicates a policy that demands more bits.

Summarizing, an RDC model of language production is thus fully specified by Eqs. 1–5, the control gain $\alpha \geq 0$ reflecting the capacity of cognitive control, the future-discount parameter $\gamma \in [0, 1]$ reflecting how much the speaker cares about future value, a need distribution over communicative goals $p(g)$ which is used to define the automatic policy in Eq. 5, and a listener model $p_L(w \mid a, s)$ which instantiates a speaker's beliefs about how a listener would interpret an utterance in terms of world states w.

Source of availability effects

The RDC production model does not have reified concepts of availability, accessibility, or activation. Nevertheless, availability-based effects emerge as a result of future discounting, which means (when $\gamma < 1$) that actions are preferred when they deliver higher value immediately. In turn, actions have high immediate value when they (1) are predictable under the automatic policy, (2) deliver high information about the speaker's target world state, or (3) set up for high value actions in the near future. For example, given the choice between including or skipping an optional complementizer, the speaker may include the complementizer if the following clause would be otherwise too unpredictable, or she may skip it if it is not as communicatively rewarding as proceeding straight into the following clause.

Study 1: Complementizer Dropping in English Relative Clauses

Speakers often have the choice to include an optional syntactic element or not in utterances such as (1):

(1) a. This is the book that everyone likes.
    b. This is the book everyone likes.

Here the optional element is the complementizer *that*; the alternation between (1-a) and (1-b) is called **complementizer dropping**.

Complementizer dropping has been a common test bed for information-theoretic models of language production. In particular, Levy & Jaeger (2007) find in a corpus study that speakers tend to drop the complementizer (i.e., *that*) when the first word of the following relative clause (i.e., *everyone likes*) is predictable in context. They attribute this finding to **Uniform Information Density** (UID): the theory that speakers include or drop the optional element in order to create a relatively even profile in terms of information content per unit time in the utterance. Under UID, if the relative clause is surprising, then the complementizer should be included in order

RC	Probability given N_1	Probability given N_2
R_1	$1/2$	$1/2$
R_2	$1/2$	$1/4$
R_3	0	$1/6$
R_4	0	$1/12$

Table 1: Need probability on worlds corresponding to utterances with nouns N_1, N_2 and relative clauses $R_1, \ldots, R_4$, used to simulate surprisal effects in relative clause complementizer dropping. Worlds corresponding to nouns N_1 and N_2 have equal total probability.

to reduce its surprisal and avoid a spike in information at the relative clause onset; on the other hand, if the relative clause is not surprising, the complementizer would create a stretch of low information density, and so should be avoided. Similar apparent UID effects have been found for other constructions (e.g., Jaeger, 2010).

These complementizer dropping results, however, are also explicable through availability-based production (Bock, 1987; V. S. Ferreira & Dell, 2000). Under this account, if the relative clause onset is highly surprising, then speakers simply have not yet formulated it yet when it comes time to produce the relative clause; they produce the complementizer to delay production and buy time to think of the right words.

Below, I show that the RDC production model predicts the effect of surprisal on complementizer dropping in the dataset of Levy & Jaeger (2007) through an availability-based mechanism.

RDC simulation

I use simulations to show that the RDC model predicts the rate of explicit complementizers to increase before high-surprisal relative clauses.

I derive predictions about complementizer dropping using a toy language in which an utterance consists of a noun, followed optionally by a complementizer, followed by a relative clause. I assume a simple listener model where each utterance x is compatible with only one target world w, and the listener places probability mass on all target worlds compatible with what has been said so far, with a small amount of probability ε distributed to other worlds:

$$p_L(w \mid x) \propto [x \text{ compatible with } w] + \varepsilon, \quad (6)$$

where $[\cdot]$ is the Iverson bracket returning 1 if the proposition inside it is true and 0 otherwise. I use $\varepsilon = 0.001$.

The toy language is designed to show effects of the surprisal of the relative clause. The language has two distinct nouns $\{N_1, N_2\}$ and four distinct relative clauses $\{R_1, \ldots, R_4\}$. The need probability on worlds w corresponding to these nouns and relative clauses is given by Table 1.

Given this setting and a controlled policy π_g, we can measure how strongly π_g favors production of the explicit complementizer as a function of the surprisal of the relative clause, by looking at the difference in (log) probability assigned to the action *that* before the high-surprisal relative clause R_2 vs. the low-surprisal relative clause R_1 after noun N_2:

$$\text{Pref}_{\text{Surprisal1}} = \ln \pi_{N_2 R_2}(\text{that} \mid N_2) - \ln \pi_{N_2 R_1}(\text{that} \mid N_2). \quad (7)$$

This value will be positive when the probability for the explicit complementizer is higher before the high-surprisal relative clause. We can also check for a surprisal preference for the *same* relative clause R_2 across different nouns N_1 and N_2:

$$\text{Pref}_{\text{Surprisal2}} = \ln \pi_{N_2 R_2}(\text{that} \mid N_2) - \ln \pi_{N_1 R_2}(\text{that} \mid N_1). \quad (8)$$

Here the comparison is done with respect to relative clause R_2 which has high surprisal after N_2 but low surprisal after N_1. This measure is also positive when the low-surprisal configuration favors production of the complementizer.

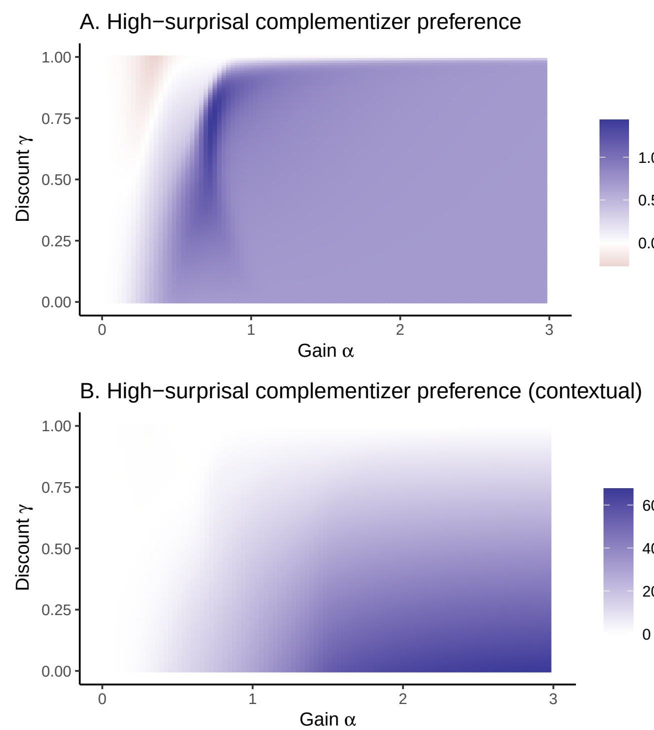

Figure 2: The RDC model-predicted bias towards explicit complementizers before high-surprisal relative clauses, as a function of control gain α and future discount γ. **A.** For different relative clauses after the same noun, as in Eq. 7. **B.** For the same relative clause after different nouns, as in Eq. 8.

Figure 2 shows the resulting surprisal preferences as a function of RDC model parameters α and γ. Across a wide range of parameters, the high-surprisal RC is predicted to be used with more complementizers, matching the UID-based prediction and results from Levy & Jaeger (2007). The effect attenuates at very high levels of the future discount parameter γ, vanishing as $\gamma \to 1$, indicating that the effect is availability-based, driven by the future-discount mechanism.

Predictor	Coefficient	95% CrI
(Intercept)	-0.88	[-1.14, -0.63]
Surprisal	0.27	[0.22, 0.33]

Table 2: Coefficients from a Bayesian logistic regression predicting presence of optional complementizer *that* before non-subject-extracted relative clauses in the dataset of Levy & Jaeger (2007). 'Surprisal' is the surprisal of the first word of the relative clause after its head noun, with or without the complementizer inserted.

Corpus Study

I verified the predicted surprisal effect in complementizer dropping using the 2904 datapoints of Levy & Jaeger (2007). Unlike the PCFG-based surprisal model used in the original paper, I calculate surprisal of the first token in attested relative clauses using GPT-2 (Radford et al., 2019). In a logistic regression predicting complementizer presence or absence as a function of surprisal of the first word of the RC (given the preceding context and the explicit complementizer), with random slopes by NP head, I find the expected positive effect of surprisal, shown in Table 2.

Discussion

The result shows that the RDC model reproduces intuitive predictions about the effect of availability on a speaker's choice to include optional elements such as complementizers. A high-surprisal RC would incur high control cost, requiring high information from cognitive control to specify the correct action. A greedy speaker can delay this cost by including the optional *that*.

Study 2: Mandarin Classifier Choice

Availability effects also appear in the choice of words that are more or less informative about following words. In particular, an apparent availability effect has been documented in Mandarin Chinese, where quantified nouns must be preceded by classifiers which may be specific or generic. For example, before the noun meaning 'computer', it is possible to use the special classifier 台 *tái* which applies to machines, or the generic classifier 个 *gè* which can be paired with almost any noun, as in (2):

(2) a. 一 台 电脑
one MACHINE computer
'one computer'
b. 一 个 电脑
one GENERIC computer
'one computer'
c. 一 只 猫
one ANIMAL cat
'one cat'
d. 一 个 猫
one GENERIC cat
'one cat'

In picture-naming experiments involving Mandarin classifier choice, Zhan & Levy (2019) found that speakers prefer the generic classifier before low-frequency nouns, and that they use the generic classifier more when under time pressure. These results support an availability-based account as follows: when it comes time to produce the classifier, the speaker may not have yet resolved which noun she will use, and thus does not know which specific classifier would be appropriate, so she uses the generic classifier instead. Since it conveys very little information in this context, the generic classifier *gè* is effectively a filled pause, with the function of delaying production of the noun (Lounsbury, 1954; Goldman-Eisler, 1957; Henderson et al., 1965; Clark & Fox Tree, 2002; Harmon & Kapatsinski, 2015, 2021). This effect is directly contrary to the predictions of UID, which would predict that people use more informative specific classifiers before surprising words in order to even out the information profile.

RDC simulation

The RDC production model predicts the Zhan & Levy (2019) pattern of results regarding the effect of frequency and time pressure on classifier choice. I conducted RDC simulations using a toy world and language presenting the Mandarin classifier domain. In the simulation, there are $N = 200$ different possible nouns, each of which is assigned randomly to one of 10 specific classifiers, and an utterance consists of a generic or specific classifier followed by a noun. The need distribution over the world states corresponding to the possible nouns n follows a power law: $p(n) \propto r_n^{-3/2}$, where r_n is the frequency rank of noun n. The listener model is the same as in the complementizer dropping study above. The future-discount factor is $\gamma = 0.9$. The effect of time pressure is simulated by varying the control gain α, with lower α corresponding to greater time pressure. This makes sense under the interpretation of α as representing the channel capacity of cognitive control: when participants must act quickly then they have access to fewer bits of information from cognitive control.

Results

The empirical data from Zhan & Levy (2019) and simulation results are shown in Figure 3A–B, in terms of the probability to produce the specific classifier (as opposed to the generic classifier) as a function of noun probability. The model reproduces the positive effect of noun probability on specific classifier use, as well as the reduced use of specific classifiers in the quick/low-gain condition.

The RDC model not only reproduces the experimentally-observed pattern, but also supports the intuition that the reason speakers do not use specific classifiers for low-frequency nouns is because they have high uncertainty about the appropriate classifier. Figure 3C shows the RDC model's probability to produce the specific classifier as a function of the policy's entropy over the set of possible specific classifiers for each of the nouns given the communicative goal. The probability to produce a specific classifier is low when the model has high uncertainty aver the appropriate specific classifier.

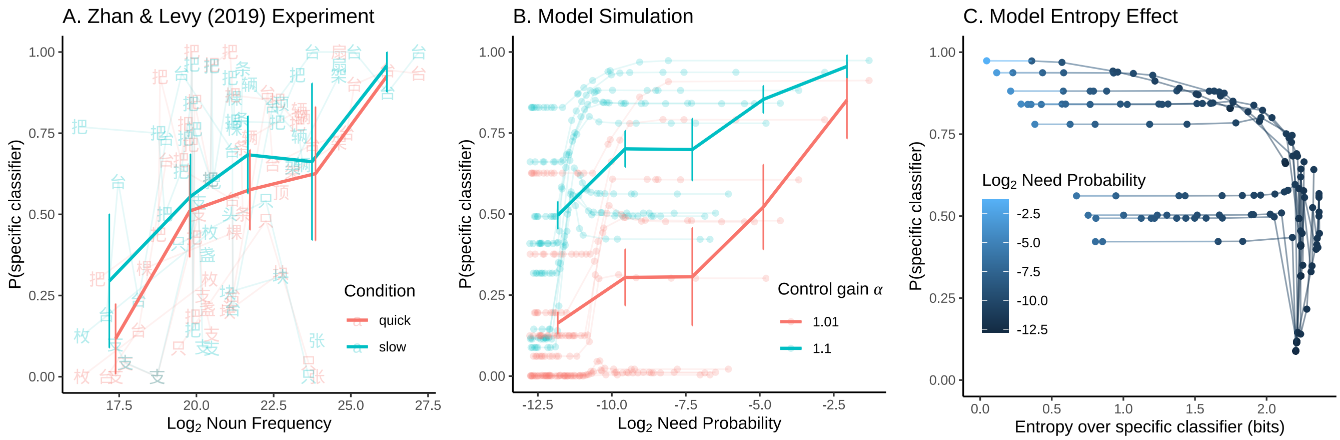

Figure 3: Experimental data and model simulations for Mandarin classifier choice. **A.** Empirical proportions of specific classifier use in a picture naming task, by noun frequency and timing condition. Solid lines show empirical means for 5 bins of noun frequency and error bars are 95% confidence intervals. Transparent lines link datapoints for individual classifiers. **B.** Model probability to use a specific classifer as a function of need probability and control gain, in a toy language with 200 Zipf-distributed nouns and 10 specific classifiers distributed randomly among them. **C.** Model probability to use the specific classifier as a function of the entropy over specific classifiers given the communicative goal, for control gain $\alpha = 1.1$.

The successful simulation of an availability effect in Mandarin classifier choice shows that the RDC production model makes successful predictions even when these are contrary to existing information-theoretic models such as UID.

Related Work

Studies of availability effects in language production are often theory-neutral in terms of the underlying causes of availability, which is usually operationalized using a combination of empirical factors such as word frequency, length, animacy, definiteness, concreteness, and others (e.g. Stallings & MacDonald, 2011; Morgan & Levy, 2016; Koranda et al., 2021). In contrast, the RDC model here provides a theoretical account of what exactly constitutes availability—high value according to Eq. 1—and why more available words are sometimes greedily selected in production: future value is uncertain and discounted, so actions with high immediate value are preferred. It remains to be seen to what extent the known availability factors such as definiteness, givenness, etc., can also be attributed to these generic mechanisms.

Production models The RDC production model used here is a computational-level model, and as such it should be hoped that it reproduces behaviors and dynamics observable in existing, more algorithmic models of language production. In some cases, components of the RDC model have direct analogues in existing models: for example the listener model p_L corresponds to the Evaluator in the Recurrent model of V. S. Ferreira (2019). More commonly, the correspondence is indirect: for example, the models here produce availability effects without explicit notions of lexical access or levels of activation, which are common in existing models (eg. Dell, 1986; Levelt, 1999; Roelofs, 2003; Dell et al., 2014).

Pragmatics models In the computational psycholinguistics literature, the RDC production model is most closely related to the Rational Speech Acts (RSA) model of pragmatic language use (Frank & Goodman, 2012; Goodman & Lassiter, 2014; Scontras et al., 2021; Cohn-Gordon et al., 2019), especially in its recent formulation in terms of rate–distortion theory (RD-RSA; Zaslavsky et al., 2021; Zhou et al., 2022). The model presented here can be seen as an incremental form of RD-RSA but with future discounting and without a pragmatic listener model: that is, the listener models used here do not depend on the speaker policy nor priors over world states as they do in RSA. The close similarity of these models suggests that a unified model of pragmatic reasoning and real-time online language production is possible.

Conclusion

I have presented an information-theoretic perspective on language production based on a combination of information theory and control theory, and shown that it can explain certain availability effects in language production. The model offers a view of language production which is formally unified with more general information-theoretic models of perception, memory, action, and neural computation that are being developed in other fields.

Acknowledgments

I thank Roger Levy for generously sharing data; Judith Degen, Michael Hahn, Greg Scontras, the MIT Computational Psycholinguistics Lab, and the UCI Quant-Lang Lab, Language Processing Group, and Center for Theoretical Behavioral Sciences for helpful comments. Code to reproduce the results here is available at `http://github.com/Futrell/rdclangprod`.

References

Andreas, J. (2022). Language models as agent models. In *Findings of the Association for Computational Linguistics: EMNLP 2022*. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.

Arumugam, D., Ho, M. K., Goodman, N. D., & Roy, B. V. (2023). Bayesian reinforcement learning with limited cognitive load. *arXiv preprint arXiv:2305.03263*.

Bhui, R., Lai, L., & Gershman, S. J. (2021). Resource-rational decision making. *Current Opinion in Behavioral Sciences*, *41*, 15–21.

Bock, J. K. (1982). Toward a cognitive psychology of syntax: Information processing contributions to sentence formulation. *Psychological Review*, *89*, 1–47.

Bock, J. K. (1987). An effect of the accessibility of word forms on sentence structures. *Journal of Memory and Language*, *26*(2), 119–137.

Bock, J. K., & Irwin, D. E. (1980). Syntactic effects of information availability in sentence production. *Journal of Verbal Learning and Verbal Behavior*, *19*(4), 467–484.

Bock, J. K., & Levelt, W. J. M. (1994). Language production: Grammatical encoding. In *Handbook of Psycholinguistics* (pp. 945–984). San Diego: Academic Press.

Bock, J. K., & Warren, R. K. (1985). Conceptual accessibility and syntactic structure in sentence formulation. *Cognition*, *21*, 47–67.

Chang, F. (2009). Learning to order words: A connectionist model of Heavy NP Shift and accessibility effects in Japanese and English. *Journal of Memory and Language*, *61*, 374–397.

Clark, H. H., & Fox Tree, J. E. (2002). Using uh and um in spontaneous speaking. *Cognition*, *84*(1), 73–111.

Cohn-Gordon, R., Goodman, N., & Potts, C. (2019). An incremental iterated response model of pragmatics. In *Proceedings of the society for computation in linguistics (SCiL) 2019* (pp. 81–90). Retrieved from `https://aclanthology.org/W19-0109` doi: 10.7275/cprc-8x17

Cover, T. M., & Thomas, J. A. (2006). *Elements of information theory*. Hoboken, NJ: John Wiley & Sons.

Dell, G. S. (1986). A spreading-activation theory of retrieval in sentence production. *Psychological Review*, *93*(3), 283.

Dell, G. S., Nozari, N., & Oppenheim, G. M. (2014). Word production: Behavioral and computational considerations. *The Oxford Handbook of Language Production*, 88–104.

Fan, J. (2014). An information theory account of cognitive control. *Frontiers in Human Neuroscience*, *8*, 680–696.

Ferreira, F., Bailey, K. G., & Ferraro, V. (2002). Good-enough representations in language comprehension. *Current Directions in Psychological Science*, *11*(1), 11–15.

Ferreira, F., & Lowder, M. W. (2016). Prediction, information structure, and good-enough language processing. *Psychology of Learning and Motivation*, *65*, 217–247.

Ferreira, F., & Patson, N. D. (2007). The 'good enough' approach to language comprehension. *Language and Linguistics Compass*, *1*(1-2), 71–83.

Ferreira, F., & Rehrig, G. (2019). Linearisation during language production: Evidence from scene meaning and saliency maps. *Language, Cognition and Neuroscience*, *34*(9), 1129–1139.

Ferreira, F., & Swets, B. (2002). How incremental is language production? Evidence from the production of utterances requiring the computation of arithmetic sums. *Journal of Memory and Language*, *46*(1), 57–84.

Ferreira, V. S. (2019). A mechanistic framework for explaining audience design in language production. *Annual Review of Psychology*, *70*(1), 29–51.

Ferreira, V. S., & Dell, G. S. (2000). Effect of Ambiguity and Lexical Availability on Syntactic and Lexical Production* 1. *Cognitive Psychology*, *40*(4), 296–340.

Ferreira, V. S., & Griffin, Z. M. (2003). Phonological influences on lexical (mis) selection. *Psychological Science*, *14*(1), 86–90.

Ferreira, V. S., & Yoshita, H. (2003). Given-new ordering effects on the production of scrambled sentences in Japanese. *Journal of Psycholinguistic Research*, *32*(6), 669–692.

Frank, M. C., & Goodman, N. D. (2012). Quantifying pragmatic inference in language games. *Science*, *336*.

Genewein, T., Leibfried, F., Grau-Moya, J., & Braun, D. A. (2015). Bounded rationality, abstraction, and hierarchical decision-making: An information-theoretic optimality principle. *Frontiers in Robotics and AI*, *2*, 27.

Gershman, S. J. (2020). Origin of perseveration in the trade-off between reward and complexity. *Cognition*, *204*, 104394.

Gershman, S. J., & Bhui, R. (2020). Rationally inattentive intertemporal choice. *Nature Communications*, *11*(1), 1–8.

Goldberg, A. E., & Ferreira, F. (2022). Good-enough language production. *Trends in Cognitive Sciences*.

Goldman-Eisler, F. (1957). Speech production and language statistics. *Nature*, *180*(4600), 1497–1497.

Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, *20*(11), 818–829.

Goodman, N. D., & Lassiter, D. (2014). Probabilistic semantics and pragmatics: Uncertainty in language and thought. *Handbook of Contemporary Semantic Theory, Wiley-Blackwell*.

Griffin, Z. M. (2001). Gaze durations during speech reflect word selection and phonological encoding. *Cognition*, *82*(1), B1–14.

Harmon, Z., & Kapatsinski, V. (2015). Studying the dynamics of lexical access using disfluencies. *Papers presented at DiSS*, 41–44.

Harmon, Z., & Kapatsinski, V. (2021). A theory of repetition and retrieval in language production. *Psychological Review*, *128*(6), 1112.

Henderson, A., Goldman-Eisler, F., & Skarbek, A. (1965). The common value of pausing time in spontaneous speech. *Quarterly Journal of Experimental Psychology*, *17*(4), 343–345.

Jaeger, T. F. (2010). Redundancy and reduction: Speakers manage syntactic information density. *Cognitive Psychology*, *61*(1), 23–62.

Kapatsinski, V. (2010). Frequency of use leads to automaticity of production: Evidence from repair in conversation. *Language and Speech*, *53*(1), 71–105.

Kemp, C., & Regier, T. (2012). Kinship categories across languages reflect general communicative principles. *Science*, *336*(6084), 1049–1054.

Kemp, C., Xu, Y., & Regier, T. (2018). Semantic typology and efficient communication. *Annual Review of Linguistics*, *4*(1), 109–128.

Koranda, M. J., Zettersten, M., & MacDonald, M. C. (2021). Good-enough production: Selecting easier words instead of more accurate ones. *Psychological Science*, 09567976221089603.

Lai, L., & Gershman, S. J. (2021). Policy compression: An information bottleneck in action selection. In *Psychology of learning and motivation* (Vol. 74, pp. 195–232). Elsevier.

Levelt, W. J. M. (1981). The speaker's linearization problem. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, *295*(1077), 305–315.

Levelt, W. J. M. (1989). *Speaking: From intention to articulation*. Cambridge, MA: MIT Press.

Levelt, W. J. M. (1999). Models of word production. *Trends in Cognitive Sciences*, *3*(6), 223–232.

Levy, R., & Jaeger, T. F. (2007). Speakers optimize information density through syntactic reduction. *Advances in neural information processing systems*, *19*, 849–856.

Lewis, R. L., Howes, A., & Singh, S. (2014). Computational rationality: Linking mechanism and behavior through bounded utility maximization. *Topics in Cognitive Science*, *6*(2), 279–311.

Lieder, F., & Griffiths, T. L. (2019). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, *43*, 1–60.

Lounsbury, F. G. (1954). Transitional probability, linguistic structure, and systems of habit-family hierarchies. In *Psycholinguistics: A survey of theory and research problems* (pp. 93–101). Waverly Press Baltimore.

Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W.H. Freeman & Company.

McDonald, J. L., Bock, J. K., & Kelly, M. H. (1993). Word and world order: Semantic, phonological, and metrical determinants of serial position. *Cognitive Psychology*, *25*, 188–230.

Mollica, F., Bacon, G., Zaslavsky, N., Xu, Y., Regier, T., & Kemp, C. (2021). The forms and meanings of grammatical markers support efficient communication. *Proceedings of the National Academy of Sciences*, *118*, e2025993118.

Morgan, E., & Levy, R. (2016). Abstract knowledge versus direct experience in processing of binomial expressions. *Cognition*, *157*, 382–402.

Ortega, P. A., & Braun, D. A. (2013). Thermodynamics as a theory of decision-making with information-processing costs. *Proceedings of the Royal Society A: Mathematical, Physical, and Engineering Sciences*, *469*, 20120683.

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. (2019). Language models are unsupervised multitask learners. *OpenAI blog*, *1*(8), 9.

Roelofs, A. (2003). Goal-referenced selection of verbal action: Modeling attentional control in the Stroop task. *Psychological Review*, *110*(1), 88.

Rubin, J., Shamir, O., & Tishby, N. (2012). Trading value and information in MDPs. In *Decision making with imperfect decision makers* (pp. 57–74). Springer.

Schach, S., Gottwald, S., & Braun, D. A. (2018). Quantifying motor task performance by bounded rational decision theory. *Frontiers in neuroscience*, *12*, 932.

Scontras, G., Tessler, M. H., & Franke, M. (2021). A practical introduction to the Rational Speech Act modeling framework. *arXiv preprint arXiv:2105.09867*.

Shannon, C. E. (1959). Coding theorems for a discrete source with a fidelity criterion. In *Ire national convention record, pt. 4* (pp. 142–163).

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, *84*(2), 127.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*, 99–118.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, *50*(3), 665–690.

Sims, C. R. (2016). Rate–distortion theory and human perception. *Cognition*, *152*, 181–198.

Sims, C. R. (2018). Efficient coding explains the universal law of generalization in human perception. *Science*, *360*(6389), 652–656.

Stallings, L. M., & MacDonald, M. C. (2011). It's not just the "heavy NP": Relative phrase length modulates the production of heavy-NP shift. *Journal of Psycholinguistic Research*, *40*(3), 177–187.

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). Cambridge, MA: MIT Press.

Tanaka, M. N., Branigan, H. P., McLean, J. F., & Pickering, M. J. (2011). Conceptual influences on word order and voice in sentence production: Evidence from Japanese. *Journal of Memory and Language*, *65*, 318–330.

Tishby, N., & Polani, D. (2011). Information theory of decisions and actions. In *Perception–action cycle* (pp. 601–636). Springer.

Todorov, E. (2009). Efficient computation of optimal actions. *Proceedings of the National Academy of Sciences*, *106*(28), 11478–11483.

Van Dijk, S. G., & Polani, D. (2013). Informational constraints-driven organization in goal-directed behavior. *Advances in Complex Systems*, *16*(2–3), 1350016.

Zaslavsky, N., Hu, J., & Levy, R. P. (2021, February). A Rate–Distortion view of human pragmatic reasoning. In *Proceedings of the society for computation in linguistics 2021* (pp. 347–348). Online: Association for Computational Linguistics. Retrieved from https://aclanthology.org/2021.scil-1.32

Zaslavsky, N., Kemp, C., Regier, T., & Tishby, N. (2018). Efficient compression in color naming and its evolution. *Proceedings of the National Academy of Sciences*, *115*(31), 7937–7942. Retrieved from http://www.pnas.org/content/115/31/7937 doi: 10.1073/pnas.1800521115

Zaslavsky, N., Kemp, C., Tishby, N., & Regier, T. (2020). Communicative need in colour naming. *Cognitive Neuropsychology*, *37*(5-6), 312–324.

Zénon, A., Solopchuk, O., & Pezzulo, G. (2019). An information-theoretic perspective on the costs of cognition. *Neuropsychologia*, *123*, 5–18.

Zhan, M., & Levy, R. P. (2019). Availability-based production predicts speakers' real-time choices of mandarin classifiers. In *Proceedings of the 41st annual conference of the cognitive science society* (pp. 1268–1274). Montreal: Cognitive Science Society.

Zhou, I., Hu, J., Levy, R., & Zaslavsky, N. (2022). Teasing apart models of pragmatics using optimal reference game design. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 44).