

resource-rational analysis – computational resources – is often avoided, in many cases reducing the approach to that of finding the optimal algorithm under task constraints – a procedure strikingly similar to rational analysis. The reason, also acknowledged by the authors themselves, is that measuring cognitive constraints has progressed slowly.

Behind this slow progress lies a fundamental difficulty in reverse-engineering the nature of the cognitive system's components. As argued by Newell (1990), each psychological experiment produces output that is the joint product of several cognitive processes. Consequently, data collected to advance our understanding of one cognitive process is marred with unexplained variance from several others, which renders determining the exact structure of the process under investigation problematic. To address this issue, Newell proposed iteratively developing and refining a model of the entire cognitive system – a *unified theory of the mind* – that provides a unified account of an ever-increasing number of psychological tasks. By jointly carving away unexplained variance from all components of the mind, such a theory would enable each subsequent experiment to ask more specific questions about the psychological process it investigates.

Newell's behest was followed by several *cognitive architectures*, the most developed among which is likely ACT-R (Anderson 2007). This architecture has incorporated the rational analysis of memory into the earlier ACT* (Anderson 1983), added the perceptual and motor processes, meticulously developed for the EPIC cognitive architecture (Meyer & Kieras 1997a; 1997b), and linked its components to regions in the brain (Anderson et al. 2016; Borst & Anderson 2017). Currently, it is able to account for behavior in hundreds of tasks in various fields, which include language learning and comprehension (Budiu & Anderson 2004; Taatgen & Anderson 2002), decision making (Marewski & Schooler 2011), driving (Salvucci & Taatgen 2010), and many others (see <http://act-r.psy.cmu.edu/publication/> for a list of publications categorized by field). This likely makes ACT-R the best source of cognitive constraints for resource-rational analysis.

Whether ACT-R is the cognitive theory of choice or not, Newell's arguments remain valid today: addressing the *identifiability problem* (Anderson 1990), the *irrelevant specification problem* (Newell 1990), or the problem of *amortization of theoretical constructs* (Newell 1990) is likely to be most successful with a unified theory of the mind that progressively incorporates multiple constraints from experiments, evolutionary arguments, and functional considerations. In my opinion, we should be devoting more efforts to develop such theories to accelerate our understanding of the mind as even the leader, ACT-R, despite its many successes, is still far from complete: It lacks theories of some fundamental components of the mind, such as emotions and tactile and other sensations, whereas many of the currently included components will likely be subjected to multiple refinements and extensions as this architecture is put to the test in new tasks.

Advancing a unified theory of the mind will naturally benefit approaches such as resource-relational analysis. Moreover, I believe that this approach might play a role in unveiling the structure of the mind similar to the role rational analysis played in developing a theory of memory. Specifically, if we maintain the assumption of optimality, we can ask under what cognitive and task constraints the empirically observed algorithms would be optimal, which could allow us to narrow down the plausible region in the space of possible computational resources. Such synergies between optimization approaches and cognitive architectures coupled with growing efforts in developing the latter will

likely lead to considerable advancements in our understanding of human cognition and behavior.

Resource-rationality beyond individual minds: the case of interactive language use

Mark Dingemans^{a,b,c} 

^aCentre for Language Studies, Radboud University, 6525 HT Nijmegen, Netherlands; ^bDonders Institute for Brain, Cognition and Behaviour, Radboud University, 6525 HR Nijmegen, Netherlands and ^cMax Planck Institute for Psycholinguistics, 6525 XD Nijmegen, Netherlands.

m.dingemans@let.ru.nl

<https://www.ru.nl/english/people/dingemans-m/>

doi:10.1017/S0140525X19001638, e9

Abstract

Resource-rational approaches offer much promise for understanding human cognition, especially if they can reach beyond the confines of individual minds. Language allows people to transcend individual resource limitations by augmenting computation and enabling distributed cognition. Interactive language use, an environment where social rational agents routinely deal with resource constraints together, offers a natural laboratory to test resource-rationality in the wild.

The target article sketches the promise of combining rational principles and cognitive constraints to understand human cognition, and singles out linguistics as one domain for work along those lines. While it touches on aspects of language rooted in individual cognition like the principle of least effort (Lestrade 2017; Zipf 1949), I want to probe the limits of the resource-rational framework by looking beyond individual minds to interactive language use, the primary ecology of human cognition (Böckler et al. 2010; Waldron & Cegala 1992). Here, under the relentless pressures of rapid-fire turn-taking (Levinson 2016) and always-on inferential processes (Enfield 2013; Goffman 1967), language provides a window onto how social rational agents deal with resource limitations in a noisy and uncertain environment.

Human language provides ample evidence of adaptation to capacity limits in social interaction (Roberts & Levinson 2017). Articulation, relatively slow compared to processes of formulation and interpretation, forms a significant bottleneck in human communication that we can bypass thanks to pragmatic inference (Levinson 2000): any content that can be left to inference need not be explicitly articulated. This puts a premium on computable and efficient heuristics for formulation and interpretation (Frank & Goodman 2012; Van Rooij et al. 2011). But as Lieder and Griffiths argue, people cope with computational complexity through heuristics as well as through habits. One way to think of language is as offering a culturally evolved store of habits – routinely deployable resources – that help outsource computation and streamline coordination (Clark 1998; Kempson et al. 2016).

A resource-rational approach may be especially promising for understanding the ubiquity of delay markers, continuers, and

repair strategies, which easily occur in up to one in five utterances (Enfield 2017; Fox Tree 1995). Whereas classic linguistic work has assumed such items are grammatically irrelevant (Chomsky 1965) or at most symptoms of trouble (Levelt 1989), resource-rationality makes it possible to account for them as optimally adaptive interactional tools (Dingemans 2017): cognitive crutches that help optimize complex rational communication under resource limitations. For instance, delay markers like “um” help word recognition by alerting the recipient that an upcoming word might need more attention (Fox Tree 2001), and repair initiators like “huh?” or “who?” allow us to gracefully recover from impending communicative trouble, something that happens, on average, at least every 84 s in conversation (Dingemans et al. 2015). With interactional tools available at every turn to review, revise, and recalibrate understanding, the dynamics of human cognition in interaction diverges radically from the one-shot models assumed in many current theories.

As a consequence, interactive language use calls into question the exclusive focus of rational analysis on individual minds. Are resource-rational approaches limited to individual cognition or could they extend to socially distributed cognition? By enabling the redistribution of attentional, cognitive, and embodied resources (Clark 2006; Hutchins 1995), interactive language use alleviates individual-bound capacity limits and can optimize performance beyond the bounds of idealized one-shot communication: an interactively scaffolded form of cognitive offloading (Risko & Gilbert 2016). The sheer frequency of the interactional tools mentioned above shows how much communication relies on this form of scaffolding (Fusaroli et al. 2017). This radically increases the error-tolerance and flexibility of cognition in interaction. It also creates opportunities to study the workings of resource-rationality in the relatively controlled environment of well-understood sequential patterns of interaction.

Communicating under noise and uncertainty requires constant cost-benefit analyses of formulating a response versus issuing a request for repair, factoring in the relative costs of different repair formats and their possible downstream consequences, all under severe time pressure and with limited cognitive resources. A systematic comparison of repair across languages and cultures shows that people everywhere deploy the repair system in efficient ways that minimize cost for the dyad as a social unit, rather than just for themselves as individual-based rational approaches might suggest (Dingemans et al. 2015): an optimal use of *distributed* cognitive resources. A similar interactive, distributed perspective is required to make sense of information-theoretical results about word meanings and ambiguity (Piantadosi et al. 2012): we can cope with ambiguity in communication only to the extent that one mind picks up the slack where the other leaves off. This means that resource-rational analysis of human cognition will need to deal not just with individual minds, but with interacting minds operating in an environment of culturally evolved metacognitive resources.

Recent work in cognitive science and cultural evolution is revisiting the Vygotskian insight that human cognition is greatly amplified by culturally evolved pieces of cognitive equipment (Bender & Beller 2014; Clark 2006; Heyes 2018). At the same time, neuroscience is increasingly concerned with understanding brain and language in the context of social interaction (Hirsch et al. 2018; Konvalinka & Roepstorff 2012; Schilbach et al. 2013). One thing that unites these approaches is their attention to how the picture of cognitive demands and resources may change radically as a result of interactionally scaffolded, socially augmented cognition. Lieder and Griffiths do not discuss cultural

evolution and social interaction as part of the environment in which heuristics and habits can be honed to become optimally adaptive, and it is unclear whether they intend resource-rational analysis to include the kinds of interactional resources discussed here: material symbols of metacognition that augment and distribute our cognitive processes. Perhaps this is the next frontier.

In sum, I applaud the call for new ways to connect psychological theory and the cognitive sciences, and would like to put forward interactive language use as a challenging yet promising domain for resource-rational approaches. As the primary ecology of human cognition, social interaction provides a rich natural laboratory for probing the leverage and limits of resource-rational analysis. Future work in this vein might focus not just on how structural aspects of language adapt to the resource limitations of individual minds, but also on how every language offers its own compendium of culturally evolved ways by which people transcend individual resource limitations and benefit from distributed cognition.

Acknowledgment. I am grateful for thought-provoking conversations with Mark Blokpoel and Iris van Rooij, and for funding from NWO grant no. 016.Vidi.185.205.

Generalization of the resource-rationality principle to neural control of goal-directed movements

Natalia Dounskaia and Yury P. Shimansky

College of Health Solutions, Arizona State University, Phoenix, AZ 85004.
 natalia.dounskaia@asu.edu
 yury.shimansky@asu.edu
<https://chs.asu.edu/natalia-dounskaia>

doi:10.1017/S0140525X19001559, e10

Abstract

We review evidence that the resource-rationality principle generalizes to human movement control. Optimization of the use of limited neurocomputational resources is described by the inclusion of the “neurocomputational cost” of sensory information processing and decision making in the optimality criterion of movement control. A resulting tendency to decrease this cost can account for various phenomena observed during goal-directed movements.

Lieder and Griffiths demonstrate a capacity of the principle of optimal use of limited computational resources (resource-rationality principle) to account for a wide variety of observations in multiple disciplines, including psychology, neuroscience, linguistics, and economics. However, they have overlooked recent developments in the field of neural control of human goal-directed movements where the explanatory power of that principle has been demonstrated. We briefly review those developments below and show how several different pieces of evidence support the resource-rationality principle.