Cognition 128 (2013) 64–75

Contents lists available at SciVerse ScienceDirect

Cognition

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

Multiply-constrained semantic search in the Remote Associates Test

Kevin A. Smith *, David E. Huber, Edward Vul

Department of Psychology, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0109, United States

ARTICLE INFO

Article history:
Received 30 May 2012
Revised 10 January 2013
Accepted 7 March 2013
Available online 12 April 2013

Keywords:
Multiply-constrained problems
Semantic retrieval
Semantic search
Problem solving
Remote Associates Test

ABSTRACT

Many important problems require consideration of multiple constraints, such as choosing a job based on salary, location, and responsibilities. We used the Remote Associates Test to study how people solve such multiply-constrained problems by asking participants to make guesses as they came to mind. We evaluated how people generated these guesses by using Latent Semantic Analysis to measure the similarity between the guesses, cues, and answers. We found that people use two systematic strategies to solve multiply-constrained problems: (a) people produce guesses primarily on the basis of just one of the three cues at a time; and (b) people adopt a local search strategy—they make new guesses based in part on their previous guesses. These results inform how people combine constraints to search through and retrieve semantic information from memory.

1. Introduction

1.1. Multiply-constrained problems

Imagine you are planning a vacation with three finicky friends. Sam wants to relax on a beach. Pat lost her passport and must stay in the United States. Alex, an amateur volcanologist, wants to visit volcanoes. What destination would satisfy everyone? People figure out that Hawaii is good choice, and regularly solve similar problems with relative ease. They combine disparate constraints to plan the best route home based on road, weather, and traffic conditions; or to prioritize work based on demands of bosses, available resources, and dependencies from other projects. These problems are all ‘multiply-constrained’: many alternatives satisfy one constraint in isolation, but the small number of acceptable solutions can only be found via all constraints.

Multiply-constrained problems have two key features: first, each of the constraints defines qualitatively different and mutually uninformative objectives, and second, there is no common currency by which to make a principled tradeoff between criteria. The first feature differentiates multiply-constrained problems from probabilistic cue combination (Ernst & Banks, 2002; Ernst & Bulthoff, 2004; Hillis, Watt, Landy, & Banks, 2004). In probabilistic cue combination each datum provides uncertain information about the same latent variable and combining the data increases certainty; for example, obtaining a more accurate estimate of the height of a ridge by combining tactile and stereoscopic percepts (Ernst & Banks, 2002). In contrast, the constraints in multiply-constrained problems provide different types of information: in the prior example, a location's distance to the beach has no bearing on its proximity to a volcano. The second criteria captures the fact that there is no information within the problem about how to weight the constraints: one cannot judge whether a location closer to the beach but further from a volcano is preferable to one with the opposite tradeoff. Thus it is possible to have multiple acceptable answers depending on how individuals decide to weight the constraints.

* Corresponding author. Address: Department of Psychology, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0109, United States. Tel.: +1 858 534 4468.

E-mail address: k2smith@ucsd.edu (K.A. Smith).

0010-0277/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.cognition.2013.03.001

Although people often solve these problems effortlessly, this apparent ease hides the computational difficulty of the task. The space of possible answers to such a problem is usually enormous (e.g., all possible vacation spots on Earth), and an exhaustive search of all possible answers is impractical. Instead, people direct their search to promising alternatives; but how? Many theories of multiply-constrained problem solving propose a two-stage process: first people *search* for a potential answer, then they *test* this candidate against all of the constraints to rate its acceptability. If the answer is considered acceptable, people will use it as a solution; otherwise, they will search for and test another potential answer. This search-test process has been proposed as a mechanism for many cognitive tasks such as hypothesis generation (Thomas, Dougherty, Sprenger, & Harbison, 2008), analogy (Forbus, Gentner, & Law, 1995), or solving word problems (Gupta, Jang, Mednick, & Huber, 2012).

In this paper we focus on the search process – how do people come up with candidate answers. Although this test process is required to identify when the search process outputs a solution, the search process can be studied separately under the assumption that, in general, people are able to recognize a good answer when it is provided (i.e., the test process does not vary greatly across different problems). We studied the search process by obtaining a sequence of guesses as people attempted to solve a multiply-constrained problem. Prior studies have typically not studied this process as it unfolds; instead they have fit models based on a single (final) answer for each problem. We hope to gain further traction on the issue by examining the search process in an 'online' fashion, under the assumption that a sequence of guesses is a subset of proposals from the true underlying search process.

We partitioned the space of human search strategies in multiply-constrained problems along two dimensions. First, how do people use the constraints to limit the pool of candidate answers? Second, how do people search through these potential answers? Here we address these questions in a novel Remote Associates Test (RAT; Mednick, 1962) paradigm by collecting sequences of responses and quantitatively evaluating the search strategies people use to explore candidate answers.

1.2. Search in the Remote Associates Test (RAT)

The goal in RAT problems is to find one word that is associated with three cues (e.g., cues: 'moon' 'dew' 'comb'; answer: 'honey'). This task illustrates key features of multiply-constrained problems: each cue indicates a different aspect of the target word ('honeymoon' relies on a different meaning of 'honey' than 'honeycomb'), and there is no principled way to trade off association to each of the three cues. Moreover, RAT problems provide a controlled environment for studying how people solve multiply-constrained problems: all constraints are of the same type (word-word relationships), and unlike many naturalistic multiply-constrained problems, RAT problems are designed to have a unique best solution.

Not only is the Remote Associates Test a controlled multiply-constrained problem, but it is also correlated with real-world problem solving ability and creativity (Mednick, 1962), so elucidating human search strategies in the RAT can inform what drives these individual differences. Moreover, RAT performance is used to measure manipulations related to creativity, such as incubation (Vul & Pashler, 2007), affect (Isen, Daubman, & Nowicki, 1987), sleep (Cai, Mednick, Harrison, Kanady, & Mednick, 2009), and performance assessment (Harkins, 2006). Although these manipulations affect RAT solution rates, the mechanisms they impact remain unknown, so characterizing search strategies in the RAT might inform how these interventions improve creativity and problem-solving.

We next review previous attempts to specify the search process employed while taking the RAT; however, we note that these studies only considered a single final answer, rather than collecting intermediate responses during the search process. Spreading activation accounts (Collins & Loftus, 1975) of the RAT proposed that the cues activate their close associates and thus jointly activate the answer, making it more likely to be produced (Bolte, Goschke, & Kuhl, 2003; Topolinski & Strack, 2008). However, these accounts did not specify the weighting scheme for the cues, the quantitative definition of 'close associate', or the process for choosing amongst equally activated words. Gupta et al. (2012) provided evidence that the search process is affected by the frequency of candidate answers, although their model assumed an equal weighting of the cues rather than testing whether this was the case. Supporting the claim that the cues are not equally weighted, Harkins (2006) found that if the answer to a RAT problem comes to mind easily when prompted by just one of the three cue words, that problem is easier to solve. However, it is possible that these easily answered RAT problems were different in other ways—for instance, the answer to these problems may have been more strongly associated with the other cues as well. Although these studies yield promising clues about how people search for an answer in the RAT, they do not fully specify the weighting scheme for the cues, and, more importantly, do not investigate dynamic changes in the weighting scheme as the search process unfolds.

In this study, we investigated how the cues act as constraints on the words produced by the search process. The number of words related to at least one of the cues is a truly vast set of words, and an unordered exhaustive search of this set would take considerable time. Instead, we suggest that the search process samples words probabilistically, such that the constraints impact the probability that a given word is considered as a potential answer. Thus we want to know how the cues combine to impact this probability: is it the case that cues act multiplicatively, meaning that candidate answers are likely to be considered only if they are related to all three of the cues, or do the cues act additively, such that a word need only be strongly related to a single cue to be considered? To explore these questions within a probabilistic sampling framework, we considered a range of stochastic search algorithms that people could be using (Russell & Norvig, 2003). *Global search* algorithms explore the search space with no sequential dependencies, such that each word is randomly and independently selected from the same set