

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

The shape of option generation in open-ended decision problems

Permalink

<https://escholarship.org/uc/item/0bx9s4m9>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Srinivasan, Gokul

Acierno, Jane

Phillips, Jonathan

Publication Date

2022

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

The shape of option generation in open-ended decision problems

Gokul Srinivasan (gokul.srinivasan.23@dartmouth.edu)

Jane Acierno (jane.c.acierno@dartmouth.edu)

Jonathan Phillips (jonathan.s.phillips@dartmouth.edu)

Program in Cognitive Science, Dartmouth College
23 N Main St. Hanover, NH 03755 USA

Abstract

There has been a small but now growing interest in studying decision making in real-world contexts where part of the problem faced by decision makers is to generate candidate options they will actually decide between. While some of this work has employed large decision spaces where options are discrete and valuation is computationally tractable (e.g., chess), very little work has focused on genuinely open-ended decision contexts that more closely mirror mundane real-world decisions. This paper leverages large language models to investigate how people generate options when facing genuinely open-ended problems. Across three experiments, we apply semantic similarity and sentiment analyses to the options that participants sequentially generate for real-world decision problems. We find that the first options generated tend to be sampled from a relatively local region of semantic space and are typically of high value. As additional options are generated, they become increasingly dissimilar and are of lower value. These patterns held both at the level of individual option generation trajectories within a given participant and at the level of individual differences across participants.

Keywords: Option generation; Natural language processing; Semantic space; Sentiment analysis; Modal cognition

Introduction

When expert chess players look at a board, they immediately have a sense for potential moves, and often the first possibility that comes to mind is, in fact, the best move (Klein, Wolf, Militello, & Zsombok, 1995). Somehow, without having explicitly evaluated the enormous number of technically possible moves at a given board state, expert chess players have the impressive ability to generate a set of good candidate moves to consider. Here, we pursue the idea that humans employ a notably similar ability in their everyday decision making. When deciding how to spend a weekend given a set of interlocking constraints, for example, people can immediately generate a small set of candidate options to consider, and those options are typically quite good (Phillips, Morris, & Cushman, 2019). While this ability may initially seem less impressive than that of expert chess players, the set of possible ways to spend a weekend is many, many orders of magnitude larger than the set of possible chess moves, suggesting that the ability of expert chess players may actually be a special case of a more general, and perhaps more impressive, ability found throughout human cognition.

Option generation has been productively studied in games like chess in large part because the set of possible

moves for a given board state is discrete and the value of each move is well-defined. In real-world decision making, neither is true, which has made the empirical study of option generation in problems of real-world complexity much more difficult. Consequently, prior research has largely proceeded by either severely restricting the set of possibilities in highly constrained experimental paradigms or explicitly asking participants to reason over a limited set of options determined by the experimenters (Kalis, Kaiser, & Mojzisch, 2013 and Smaldino & Richerson, 2012 for further discussion.) While this prior work has been foundational for understanding reasoning and decision making about constrained sets of options, option generation in open-ended problems of real-world complexity remains under-explored relative to its centrality in everyday decision making.

Prior work on option generation. A few recent studies have explored increasingly open-ended decision problems, e.g., “What food would you most like to have for dinner?” (Morris, Phillips, Huang, & Cushman, 2021; Zhang et al., 2021). While the set of relevant options is still limited (e.g., only food eaten at dinner), exhaustive search through the set of options becomes computationally impractical. Such partially open-ended generation tasks have generated a number of important insights into the nature of option generation. Across these studies, a notably similar picture has emerged: participants generate a relatively small set of options for explicit evaluation, and the process of option generation is biased towards options that are historically valuable, likely, and semantically accessible (Morris et al., 2021; Zhang et al., 2021; Bear, Bensinger, Jara-Ettinger, Knobe, & Cushman, 2020). In line with predictions from Johnson and Raab (2003), the possibilities that come to mind *first* often rank most highly in objective and subjective value (Morris et al., 2021). Thus, in cases where there is agreement on what the highest value options are (or the semantic accessibility of options), there is a corresponding alignment on the options that first come to mind (Klein et al., 1995). Moreover, despite relatively frugal option-sampling procedures due to temporal and computational limitations, participants tend to generate consistently valuable options, echoing results from Vul, Goodman, Griffiths, and Tenenbaum (2014).

The present research. Here, we ask whether these patterns extend to more open-ended decision problems, where participants are given a series of different background decision contexts and asked to sequentially generate possible actions that could be taken. For example, one such context participants were given was:

Your significant other has recently fallen ill and needs an expensive medication that is not covered by your medical insurance. You do not have the money needed to purchase the expensive prescription medication, but you know that it is vital for them to have it if they are going to recover. In this situation, what are some things you could do?

To illustrate further, another context instead involved going on a hiking trip in Arizona where your friend slipped and gets her arm trapped in a crevice, without service and the ability to call 911. Yet another involved going to a concert with friends, but, upon reaching the concert, discovering that one of your friends had forgotten his ticket. Eighteen different decision contexts were used in across our studies.

In such cases, the set of options is clearly unbounded and ill-defined. We may naturally think of asking a friend for money, but nothing is stopping us from also considering the possibility of trying to surf on a large cheese grater across the Moab desert. Accordingly, an obvious and immediate challenge facing such an approach is how to objectively characterize the options that participants generate either relative to the set of all possible options or relative to each other.

We propose that progress can be made on this problem by leveraging recent advances in large language models. Specifically, because large language models like BERT are trained on billions of language examples, they can be used to locate the possibilities participants generate relative to the entire corpus of sentences in the training data. That is, words, phrases, and sentences can all be given numeric coordinates that represent their location in semantic space. Thus, we can think of the vector representation assigned to a given option a participant generates as occupying a point within the parameter space the model used to capture the entire corpus of sentences in the training data. Thus, large language models provide a tool we can leverage to objectively characterize the option generated even when it is a constituent of an unbounded set of options. Moreover, because we can iteratively do this for each option a participant generates, we can characterize the shape or trajectory of option generation both within a given participant's responses (how each participant explores the space of options) and across participants' responses (how people collectively search for solutions to open-ended problems). In addition, we employ a similar approach by investigating the sentiment of each option generation.

Methods

Study design

Three separate experiments were conducted in which a total of 477 participants ($N_{study1} = 197$, $N_{study2} = 178$, $n_{study3} = 102$) were recruited from Prolific (Study 1 and 2), and Prolific (Study 3) ($M_{age} = 41.2$; $SD_{age} = 12.4$; 56% women). All studies employed a similar design. Participants read a number of background contexts, ranging from 8 (Study 3) to 10 (Study 1 and 2), which each described a unique open-ended decision problem. The order of presentation for the decision-contexts was randomized.

After reading the background context, participants across studies were asked to sequentially generate a number of options that could be pursued given the problem faced (6 in Study 1 and 2; 8 in Study 3). Subsequently, participants in Study 2 and 3, were re-presented with the options they previously generated and asked to provide subjective ratings of them. In Study 2, participants rated the extent to which they believed each was a "good" option on a scale from 0 ('worst') to 100 ('best'); in Study 3, participants provided ratings of the extent to which they agreed the option they generated was rational, moral, normal, and probable on a scale from 1 ('disagree') to 7 ('agree').¹

Analysis approach

Below, we describe the two key dependent measures we use to characterize the options that participants generated. For both of these, we go on to ask how they vary within and across participants as they explore options in open-ended decision tasks. Importantly, both measures leveraged BERT-based language models that utilized novel, transformer architectures to learn the contextual relations between words (or sub-words) in a text. Whereas previous models were constrained by directional properties (reading the text input sequentially left-to-right or right-to-left), the transformer encoder reads the entire sequence of words concurrently.

Option Sentiment We utilized the fine-tuned BERT Base uncased model described in Jones and Wijaya (2021), as it achieves state-of-the-art (SOTA) or nearly SOTA results on various text classification tasks. For English training and development data, they sampled 50K positive and 50K negative tweets from the automatically annotated sentiment corpus described in Go, Bhayani, and Huang (2009) and used 90K tweets for training and the rest for development. See Jones and Wijaya (2021) for more extensive documentation on the training process.

¹Study 2 and 3 recorded the time of each sequential generation. Participants in Studies 1 and 2 were also asked to complete a range of individual difference measures, including the *Beck Depression Inventory* (BDI), the *Beck Anxiety Inventory* (BAI), the *Positive and Negative Affect Schedule* (PANAS), and the *Openness* subscale from the *Big Five Inventory*. Response time and individual difference measures are not discussed further because of space constraints.

Option Similarity We utilised the paraphrase-multilingual-MiniLM-L12-v2 model produced by Reimers and Gurevych (2019) to create a 384 dimensional dense vector embedding from sentences and paragraphs. This model, a subset of their Sentence-BERT (SBERT) model, is a modification of the BERT network using siamese and triplet networks that is able to derive semantically meaningful sentence embeddings. SBERT adds a pooling operation to the output of BERT/RoBERTa to derive a fixed sized sentence embedding. In order to fine-tune BERT/RoBERTa, they create siamese and triplet networks to update weights so that the sentence embeddings produced are semantically meaningful and can be compared with cosine-similarity. Refer to Reimers and Gurevych (2019) for further description concerning model structure and training.

Results

Prior work on option generation demonstrated that participants generate a notably similar set of high value options (Klein et al., 1995). We begin by asking whether these patterns extend to option generation in open-ended decision contexts where options space is unbounded and ill defined. Using word embeddings, we first ask whether the options participants initially generate are more similar to each other than the options they subsequently generate, which would suggest a similar alignment on an initial set of options.

Similarity of options as a function of number of options generated

Approach. To ask whether there was more semantic similarity in the first options generated by participants than later options generated, we devised a measure of semantic space similarity to capture the extent to which respondents explored similar regions of a shared semantic space (across scenarios). Formally, this likeness is measured by creating as series of localizing vectors for each participant:

$$= \frac{1}{n} \left[\sum_{i=0}^n v_{1,n} \quad \sum_{i=0}^n v_{2,n} \quad \dots \quad \sum_{i=0}^n v_{j,n} \right] \quad (1)$$

Where n is equal to the number of vignettes in the trial, v_l is equal to the l^{th} index of the embedding vector, and j is the length of the embedding vector. Concretely, the localizing vector for the 1st generation is the average vector derived from all 1st generations the participant created throughout the study. A localizing vector represents the approximate location of a participants i^{th} generations. Thus, participants will have 6 localizing vectors in Study 1 and Study 2, and 8 in Study 3.

The extent to which participants occupy distal regions of semantic space, then, can be represented as the distance between their localizing vectors. Thus, semantic space dissimilarity measures the distance between every participant’s localizing vector for a given generation number, enumerating over all $\binom{n}{2}$ combinations present in E_i , the set of

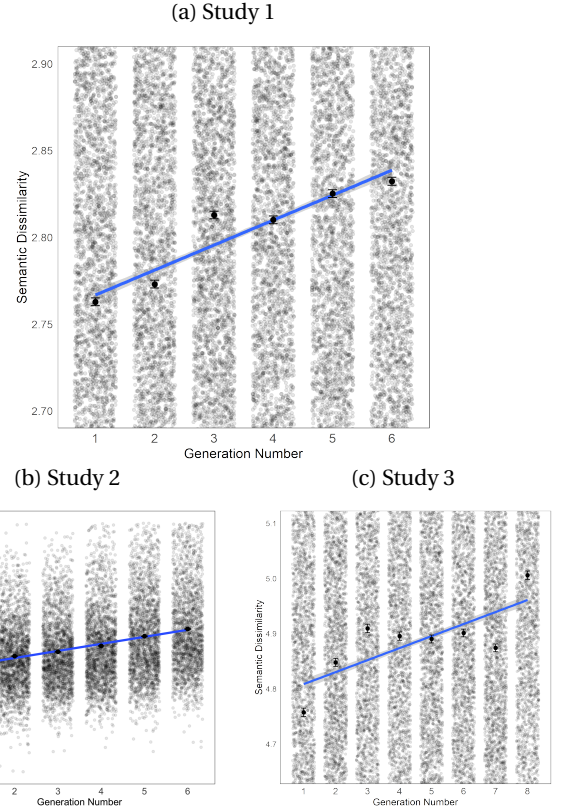


Figure 1: Depiction of the relationship between the semantic dissimilarity of the options generated as a function of generation number for all three studies: Study 1 (Fig. 1a), Study 2 (Fig. 1b), and Study 3 (Fig. 1c). Larger dots depict mean dissimilarity for each generation.

all localizing vectors for the i^{th} generation. Specifically,

$$\text{Semantic space dissimilarity}_i = \sum_{l,k \in E_i: l \neq k} \text{distance}(l, k) \quad (2)$$

Result. We found that semantic space dissimilarity varied positively with generation number across all three datasets: Study 1 ($\rho = .912, p = .011$), Study 2 ($\rho = .980, p = .001$), Study 3 ($\rho = .819, p = .013$), see Fig. 1. In other words, while participants exhibited a higher degree of semantic convergence early on—exploring semantic space in a more similar way—they exhibited a lower degree of semantic convergence later on. We conceive of semantically convergent possibilities as generations that encode similar (or identical) ideas.

Semantic exploration over generations

In the prior analyses, we found that for a given open-ended decision context participants were more aligned on the first options that came to mind but diverged from each other as they explored additional options. A separate question concerns whether there is a particular method or shape to the way in which a given participant explores options *across de-*

cision contexts. In other words, we want to know whether participants gravitate toward similar options as they explore solutions to different kinds of open-ended problems.

Approach. To explore this possibility, we created a new metric called semantic exploration, which refers to intra-participant exploration across serial positions. In other words, semantic exploration for the i^{th} generation involves summing the range of semantic space traversed across all i^{th} generations which occurred across different decision contexts. This relationship is captured mathematically as follows, where $E_{i,j}$ is the set of all i^{th} generations for the j^{th} participant, and l and k are embedding vectors:

$$\text{exploration}_{i,j} = \sum_{l,k \in E_{i,j}: l \neq k} \text{distance}(l,k) \quad (3)$$

Thus, the total exploration for the i^{th} generation is the sum of individual explorations across all P participants:

$$\text{Total exploration}_i = \sum_{j \in P} E_{i,j} \quad (4)$$

Result. Semantic exploration was negatively correlated with generation number across Study 1 ($\rho = -.860, p = .028$) and Study 3 ($\rho = -.712, p = .047$) see Fig. 2.² In other words, for a given participant, the first options that came to mind across different contexts tended to be unrelated to one another. However, as participants began to produce successive options across different contexts, the options they generated became increasingly similar to one another. For example, in the trapped arm decision-context mentioned previously, a participant provided the following generations (in-order): "cut her arm off", "hunker down", "do nothing", "ask friends for help", "daydream", and "nothing more can be done". In another, separate decision context, the same participant responded with the following generations: "pay the difference myself", "tell the manager what happened", "ask the other staff to help cover the difference", "call the police", "do nothing", and "ask the manager what to do". As can be seen in this example, the participant's later generations shared more similarity than their earlier generations. Intuitively, one can think about this as participants moving toward more generalizable (and thus more similar) solutions to different kinds of problems.

The subjective value of options generated

Following prior research, we previously demonstrated a notable alignment in the similarity of the options that first come to participants' minds for a given open-ended decision context. Existing work additionally suggests that the options that first come to mind tend to be high in value (Klein et al., 1995; Morris et al., 2021) and that subsequently generated options actually decrease in value (Johnson & Raab, 2003). We next asked whether this was also the case for option generation in open-ended decision contexts.

²We were unable to analyze the data from Study 2 due to a technical error.

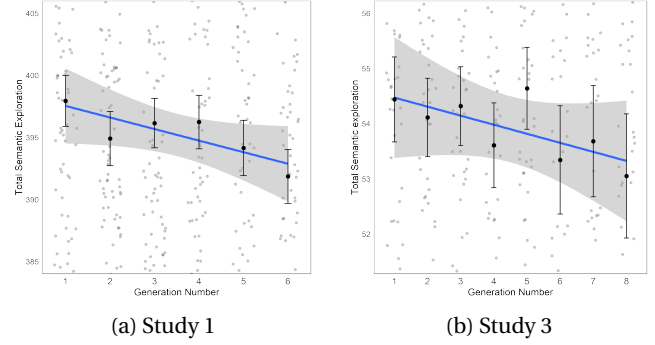


Figure 2: Depiction of the relationship between the semantic exploration within a participant/context as a function of generation number for Study 1 (Fig. 2a) and Study 3 (Fig. 2b). Larger dots depict mean total semantic exploration for each generation.

Approach. We used a diversity of ways of measuring value, which all revealed similar patterns. In Study 2, participants were asked to rate each of the actions they generated in terms of their "goodness", and in Study 3, participants were asked to rate each of the actions they generated in terms of whether they would be "rational" to do. We used both of these as measures of the *subjective* value of the options generated for each generation number. Hence, the average reflection score for the i^{th} possibility number was calculated by averaging subjective ratings across all i^{th} generations for every participant, where E_i is that set of all i^{th} generations (e.g., 150 participants over 10 vignettes = 1500 i^{th} generations = n).

$$\text{Average rating}_i = \frac{1}{n} \sum_{g \in E_i} \text{rating}(g) \quad (5)$$

Result. We found that generation number was strongly negatively correlated with both subjective ratings of "goodness" in Study 2 ($\rho = .992, p < .001$) and "rationality" in Study 3 ($\rho = .949, p < .001$)³ In other words, we replicate the finding in finding in prior work that the options that come to mind first tend to be high in subjective value, and that subsequent actions decrease in value.

Sentiment analysis as an objective estimate of the value of an option

We next ask whether we can move beyond participants' subjective ratings of the value of options which may be subject to so-called "self-serving biases" (Schlenker, 1980), and demonstrate a similar pattern for a less subjective measure of value. To do so, we ask whether sentiment analysis of the options generated would replicate the inverse relationship between an options' value and generation number.

³Here we use rationality as the most obvious way of estimating value in decision-making contexts. Similar patterns emerged for the other ratings

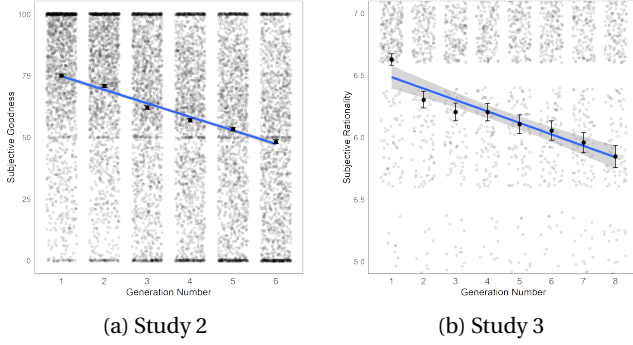


Figure 3: Depiction of the relationship between the subjective assessments of the value of options and generation number for two studies: Study 2 (Fig. 3a), Study 3 (Fig. 3b). Larger dots depict mean ratings for each generation.

Approach. Average sentiment for the i^{th} possibility number was calculated by predicting the sentiment for all i^{th} generations across every participant, where E_i is that set of all i^{th} generations.

$$\text{Average sentiment}_i = \frac{1}{n} \sum_{g \in E_i} \text{sentiment}(g) \quad (6)$$

Result. We found that the average sentiment value was strongly negatively correlated with generation number in all three datasets: Study 1 ($\rho = -.824, p = .044$), Study 2 ($\rho = -.954, p = .003$), Study 3 ($\rho = -.789, p = .020$), see Fig. 4. In short, employing sentiment analysis as a less subjective estimate of an option’s value, we find that the first options that come to mind for open-ended decisions tend to be the highest in value, and that subsequent options seem to linearly decrease in value.⁴

Individual variation in exploration and sentiment

Previously, we found that the first options that participants generate tend to both be high in value and similar, but as participants explored additional options, they became more dissimilar and lower in value. We next wanted to know whether this relationship between exploration and value held at the level of individual differences in participants. That is, we asked whether it is the case that participants who explored more also generated lower value options.

Approach. Mathematically, we investigate this intuition by defining semantic exploration as follows, where V is the set of all vignettes, v is a particular vignette, and i and j are option generations within vignette v :

$$\text{Semantic exploration} = \frac{1}{|V|} \sum_{v \in V} \sum_{i, j \in v: i \neq j} \text{distance}(i, j) \quad (7)$$

⁴It may also be worth noting that we can estimate the relationship between the sentiment score from the language model and participants subjective ratings of “goodness” in Study 2. We find that the two are clearly positively related, $\rho = .278, p < .001$

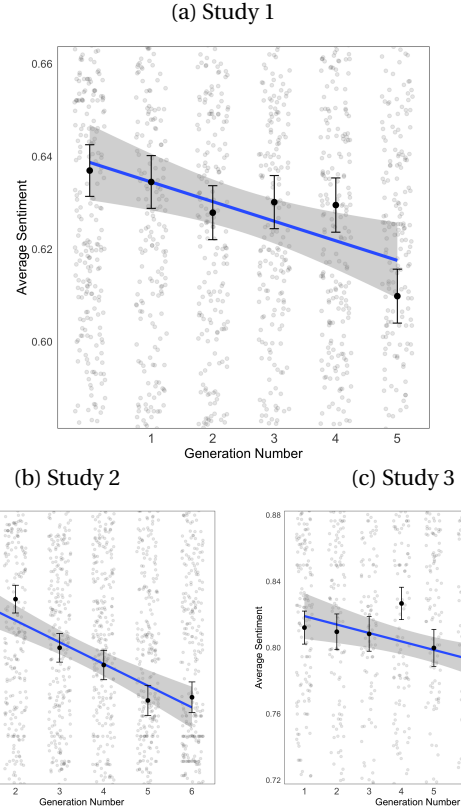


Figure 4: Depiction of the relationship between the average sentiment of options generated and generation number for all three studies: Study 1 (Fig. 4a), Study 2 (Fig. 4b), and Study 3 (Fig. 4c). Larger dots depict mean sentiment for each generation.

Further, we defined average sentiment to be the average across all of the options the participant generated within their trial.

Result. We found that individual differences in semantic exploration are inversely related to average sentiment across two of the three datasets: Study 1 ($\rho = -.468, p < .001$), Study 2 ($\rho = -.246, p = .001$), Study 3 ($\rho = .173, p = .087$), see Fig. 5. That is, we found suggestive evidence that participants who tended to explore more diverse regions of semantic space also tended to generate options with lower value as estimated by sentiment analysis. This finding however, may be sensitive to the particular decision problems being solved, as we did not see this relationship in Study 3. Moreover, these results did not replicate when substituting average subjective ratings for sentiment analysis.

Discussion

Large language models were used to analyze the structure of participants’ option generation across 3 studies involving open-ended decision contexts. As predicted by results detailed in Bear et al. (2020), Zhang et al. (2021), Klein et al. (1995), and Morris et al. (2021), participants tended to

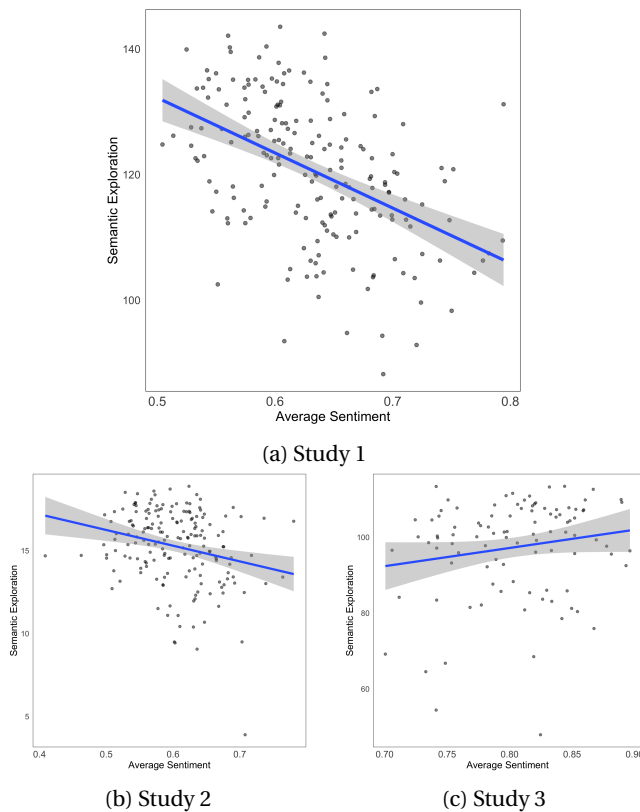


Figure 5: Depiction of the relationship between the average semantic exploration per participant across contexts and the average sentiment of the options generated by that participant for all three studies: Study 1 (Fig. 5a), Study 2 (Fig. 5b), and Study 3 (Fig. 5c).

explore relatively similar regions of semantic space in the early stages of option generation. We see this as confirmation that the process of option generation is biased towards options that are historically valuable, likely, and semantically accessible. The best possibilities will inhabit well-defined, but concentrated, pockets of semantic space, and as participants reliably generate these possibilities first, they exhibit a larger degree of semantic similarity early on. Moreover, as there are necessarily more ways of generating bad options than good options, all else held equal, greater exploration should lead to lower average sentiment scores. The inverse relationship between semantic exploration and average sentiment further confirms this idea.

Several theories also predict that the quality of an option generation will be inversely related to its generation number (Johnson & Raab, 2003; Morris et al., 2021). While our findings confirm this prediction, they stand, at first glance, in opposition to proposals put forth by Lieder, Griffiths, and Hsu (2018), in which option-generation mechanisms should over-represent options with both extremely positive and extremely *negative* qualities. However, respondents in Lieder et al. (2018) generated options under significant uncertainty. In our studies, by contrast, the participant’s task

was to select actions that they themselves could do. Because participants had full control over their actions, there was no need to over-represent negative extremes, as there was no reason to fear that they may select options of low value.

In addition, the positive relationship between semantic space dissimilarity and generation number confirms another prediction espoused by Johnson and Raab (2003)—that successive option generations should increasingly diverge from the original generation. On their view, option generation is composed of distinct construction and retrieval systems driven by spreading activation. Accordingly, they argue that option space is traversed in a Dijkstra-like fashion based on the strength of the semantic connection between options (Johnson & Raab, 2003).

The inverse relationship between total semantic exploration and generation number suggests that the first options that came to mind in different contexts tended range across diverse regions of semantic space. However, as participants began to generate more options for that context, the options they generated became increasingly related to other options across different contexts. Prima facie, there is no clear reason that later option generations for disparate decision problems should converge on shared regions of semantic space—yet, this is what we observe. In combination with our finding that later options also decline in both subjective and objective value, a humorous picture arises: each participant, when they are solving a range of different problems, tend to head towards a semantically similar, low value part of semantic space. One possibility is that this phenomenon is a function of decision fatigue. As participants experience increasing fatigue, they may begin to default to a relatively more domain-general set of options which happen to be less high in general value when implemented in specific contexts. Regardless, these results represent a fruitful avenue for further study.

In closing, we want to emphasize that use of language models provides a fecund analytic paradigm for delineating the manner in which people solve open-ended decision problems. While our investigation focused on employing a few robust and well-understood semantic techniques, there exist a much wider range of tools this work leaves remarkably underexplored (e.g., unsupervised clustering algorithms or supervised learning algorithms). Future work will certainly benefit from these and other novel tools that will no doubt emerge given the rapid pace of advancement within the domain of natural language processing.

Acknowledgments

We would like to thank Camden Parker and Yuri Oh for their assistance in conceptualizing the study.

References

Bear, A., Bensinger, S., Jara-Ettinger, J., Knobe, J., & Cushman, F. (2020). What comes to mind? *Cognition*, 194, 104057.

- Go, A., Bhayani, R., & Huang, L. (2009). Twitter sentiment classification using distant supervision. *Processing*, 1–6.
- Johnson, J. G., & Raab, M. (2003). Take the first: Option-generation and resulting choices. *Organizational behavior and human decision processes*, 91(2), 215–229.
- Jones, A., & Wijaya, D. T. (2021). *Sentiment-based candidate selection for nmt*.
- Kalis, A., Kaiser, S., & Mojzisch, A. (2013). Why we should talk about option generation in decision-making research. *Frontiers in psychology*, 4, 555.
- Klein, G., Wolf, S., Militello, L., & Zsombok, C. (1995). Characteristics of skilled option generation in chess. *Organizational behavior and human decision processes*, 62(1), 63–69.
- Lieder, F., Griffiths, T. L., & Hsu, M. (2018). Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psychological Review*, 125(1), 1.
- Morris, A., Phillips, J., Huang, K., & Cushman, F. (2021). Generating options and choosing between them depend on distinct forms of value representation. *Psychological science*, 32(11), 1731–1746.
- Phillips, J., Morris, A., & Cushman, F. (2019). How we know what not to think. *Trends in cognitive sciences*, 23(12), 1026–1040.
- Reimers, N., & Gurevych, I. (2019). Sentence-bert: Sentence embeddings using siamese bert-networks.
- Schlenker, B. R. (1980). *Impression management* (Vol. 222). Monterey, CA: Brooks/Cole.
- Smaldino, P. E., & Richerson, P. J. (2012). The origins of options. *Frontiers in Neuroscience*, 6, 50.
- Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? optimal decisions from very few samples. *Cognitive Science*, 38(4), 599–637.
- Zhang, Z., Wang, S., Good, M., Hristova, S., Kayser, A. S., & Hsu, M. (2021). Retrieval-constrained valuation: Toward prediction of open-ended decisions. *Proceedings of the National Academy of Sciences*, 118(20).