

Modeling Idea Generation Sequences Using Hidden Markov Models

Hao-Chuan Wang (haochuan@cs.cmu.edu)

Language Technologies Institute, Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA 15213 USA

Abstract

The paper presents computational models for investigating cognitive and socio-cognitive processes in the task of idea generation. We employed Hidden Markov Model (HMM) and its variants to model memory structures and processes in idea generation that have been theorized descriptively but not implemented and tested computationally. Existing algorithms and methods for HMM provide a new path to examine these invisible structures and processes, and the influence of social interactions in brainstorming groups.

Keywords: Idea Generation; Cognitive Modeling; Hidden Markov Model

Introduction

In the past decades, laboratory experiments on group idea generation have demonstrated many interesting phenomena regarding human creativity and how group may affect the mind. One persistent finding is the phenomenon called process loss that real brainstorming groups characterized by information sharing and social interactions actually generate significantly fewer unique ideas than nominal groups in which individuals work alone and their ideas are pooled together afterward (Diehl & Stroebe, 1987; Brown et al., 1998; Nijstad & Stroebe, 2006; Wang & Rosé, 2007). The finding is contrary to the popular belief that working in groups is better than alone or social interactions would be beneficial for creative tasks.

In search for theoretical explanations for process losses in group brainstorming or group idea generation, historically, laboratory experiments played the major role in testing plausible explanations. Classic explanations are mostly at the social level, such as production blocking due to turn taking, social loafing (“free riding”), and evaluation apprehension. Controlled experiments were able to test hypotheses about these social factors and established causation (Diehl & Stroebe, 1987). Recently, resurging interests on the research of idea generation have shifted the focus to the cognitive aspects, in particular, how interactions between two or more individual cognitive systems lead to observed group outcome. The socio-cognitive approach explain the target group phenomenon as a consequence of interactions between individuals, and therefore one’s memory retrieval and verbalization is affected by partners’ behaviors. However, studying cognitive processes in social contexts through behavioral studies could be expensive and methodologically constrained. For example, it could be difficult to manipulate participants’ behaviors for testing cognitive explanations and at the same time ensure the authenticity of social interactions.

Computational modeling of human cognition is potentially useful for studying the cognitive process of

individual idea generation and the socio-cognitive process regarding how social interactions influence individuals’ ideation. It is plausible to model idea generation as a Markov stochastic process based on the observation that the generation of a later idea may be influenced by ideas that have been triggered. Previous work has modeled idea generation as Markov chains (Brown et al., 1998). The modeling formalism was effective and simple, but did not show adequate modularization and abstraction that modern cognitive theories have generally agreed upon.

The objective of the work is to address the limitation of prior work by introducing Hidden Markov Model (HMM) (Rabiner, 1989) and its variants to model idea generation. Hidden states in HMMs are leveraged as a computational device to model non-observable memory retrieval processes during idea generation have been theorized by Nijstad & Stroebe (2006) but not yet fully explored through computational modeling.

This paper consists of three components. First, we review the state-of-the-art research on cognitive modeling for idea generation. We then train HMMs with empirical data to model individual idea generation. To demonstrate the utility, we perform two analyses made possible by HMM, including (1) examining the structures of mental images and (2) examining the transitions of mental states behind idea generation sequences. Lastly, we model interactive idea generation in groups using Input-Output HMM (IO-HMM) (Bengio & Frasconi, 1996). As the third analysis of the paper, we compare two plausible modeling structures of IO-HMM and discuss possible ways that interactive communication may affect idea generation.

Cognitive Modeling for Idea Generation

Idea generation can be viewed as a process of searching for ideas in the long-term memory (LTM), in which concepts are assumed to store in network structures, namely the associative memory (Nijstad & Stroebe, 2006). Thinking of one concept may activate another concept through encoded associations. The generation of a later idea thus may possess dependency on previous ideation. For example, when thinking ideas for the problem, “how to improve university parking”, the idea of encouraging students to walk is more likely to co-occur with the idea of encouraging students to ride bicycles but probably not the idea of increasing the parking fee. From the view of cognitive modeling, the possible local dependency between consecutive ideas in idea generation sequences is one of the salient features that a cognitive modeling framework strives to capture and describe.

Two existing cognitive modeling frameworks for idea generation are both based on the assumption that

spreading activation in associative memory drives the production of ideas and related behavioral phenomena. (Brown et al., 1998; Nijstad & Stroebe, 2006). Brown et al. (1998) modeled idea generation as a Markov chain consisting of transitional probabilities for a set of idea categories without assuming hidden states. In this model, every specific idea is assumed to have only one categorical label. To generate a specific idea requires accessing the predefined category that the idea is associated with. At each time stamp, one may either generate a within-category idea that is under the same category as the previous idea, or transit to another category and generate an idea from the new one. This modeling approach is easy to handle, and was shown useful in simulating idea generation at individual and group levels. However, despite its simplicity, as Nijstad & Stroebe (2006) argued, employing only a probabilistic transitional matrix to describe idea generation behaviors failed to acknowledge established cognitive theories that possess higher degrees of complexity and modularity. Also, the assumption that every idea associates with only one category is debatable

Nijstad & Stroebe (2006) proposed an alternative modeling approach as an attempt to better incorporate cognitive theories. Based on a well-established theory of associative memory (Raaijmakers & Shiffrin, 1981), Nijstad & Stroebe’s SIAM (Search for Ideas in Associative Memory) model assumed that memory retrieval primarily consists of two stages. First, an image r containing a central concept, related concepts, and properties associated with concepts is retrieved from LTM to the working memory (WM). Ideas are subsequently generated from the activated image r . The SIAM model explicitly draws parallels between idea generation and free recall to avoid overly specializing idea generation and to explain related phenomena based on general cognitive theories. However, the SIAM model is at best a procedural flowchart but not a fully parameterized computational model. The model is useful for hypothesis formation and testing, which emphasizes on the direction of prediction but not precise quantification. The model cannot produce simulations with sufficient details, such as predicting which idea is more/less likely to occur given the current mental state.

To bridge the gap between Brown et al.’s pragmatic approach of Markov chain modeling and Nijstad & Stroebe’s theory-driven approach, we propose to use HMM as a computational device of modeling that captures observable ideation outputs and theorized, non-observable activation of memory sub-structures (e.g., images in SIAM). The HMM approach is plausible based on these discussions. Recent success in using HMM to model various cognitive processes, such as eye movements (Feng, 2006) and cognitive load in teamwork (Fan et al., 2007), also motivates this approach.

To model idea generation, we assume that semantic memory is partitioned into sub-structures called mental images and can be represented as hidden states in HMM. Ideas are output symbols emitted from the hidden states. The observed semantic dependency between consecutive ideas can be explained by having ideas generated from the same or different states. Theorized processes like

Table 1: Codes for the “Debris Flow Hazard” idea generation task

1	Steep mountain slope	12	Ineffective water drainage
2	Fragile geological condition	13	Illegal sand mining
3	Deforestation	14	Excessive earth and water pressure
4	Growing shallow root plants	15	Ineffective flood prevention
5	Intense rainfall	16	People living in unsafe areas
6	Earthquake	17	Lack of information and education
7	Excavating slope toe	18	Lack of awareness
8	Excessive urban development	19	Bad policy and/or execution
9	Inadequate soil and water conservation	20	[Unrelated idea]
10	Building houses on the steep slope	21	[Proposal of solution]
11	Improperly dumping waste soil	22	[End of ideation]

abandoning the current image and activating a new one or exhausting ideas from a few images may be captured by probabilistic transitions between hidden states. In this way, we may employ HMM to model individual idea generation in which external cues are absent. Given this foundation, we further use IO-HMM to account for the influence of overhearing others’ ideas on individuals in brainstorming groups.

Modeling Individual Idea Generation

Dataset and Preprocessing

We trained HMM for individual idea generation by using a dataset consisting of 226 Taiwanese high school students’ idea generation sequences on a brainstorming topic in the domain of Earth Sciences. These students were asked to answer the question, “*What are possible factors for the happening of Debris Flow Hazards (DFH)?*” DFHs are natural hazards frequently occurred in Taiwan. Students worked individually to generate a list of ideas on the open-ended question in 30 minutes.

Their responses were coded manually as one of the 19 valid ideas identified by domain experts. For responses that are not identified as any of the 19 valid ideas, the code “unrelated idea” was assigned. For responses that were related but talking about possible solutions for the hazard, the code “solution” was assigned (the task asked for underlying factors, not solutions). At the end of each idea sequence, we appended another unique code, “end of ideation”, to indicate running out of ideas. Therefore a total of 22 codes were used to code each idea sequence. Table 1 shows a brief summary of the coding scheme. Two independent coders coded the dataset as reported in Wang, Chang & Li (in press). The inter-coder agreement achieved the level of Kappa = .77, which was satisfactory. Among the




Figure 1: Likelihood-based cross-validation for choosing the number of hidden states

226 sequences (one from each student), the longest sequence consisted of 16 codes, and the shortest consisted of 3 codes.

Learning Parameters of HMM

We assumed that individual idea generation can be modeled by HMMs with n hypothetical hidden states, where $3 \leq n \leq 10$. Since hidden states are not observable, deciding the best n is a critical step. We used leave-one-out cross-validation (CV) (i.e., 226 folds CV) to avoid over-fitting and select the most generalizable model. At each fold, we held one sequence as the testing data and trained an n -state HMM model using the rest 225 sequences. We tested the likelihoods of the testing data given the trained model. The summation of log-likelihoods derived through cross-validation was employed as the goodness of fit index. We trained HMMs by employing an implementation of the Baum-Welch algorithm included in the HMM toolbox for Matlab (Murphy, 1998). When estimating emission probabilities, we utilized the popular add-one smoothing technique to add one imaginary count to each observable codes.

Figure 1 shows the goodness-of-fit scores (summations of log-likelihoods computing from leave-one-out CV) associated with each of the n -state HMMs. From the comparison, HMM with 8 hidden states was the optimal. We decided to use the 8-state model in following analyses.

Analysis 1: Structure of Mental Images

Figure 2 shows the distribution of emission probabilities in the 8-state model. The distributions reveal how ideas are clustered around the mental images in HMM. It makes sense that code 20 (unrelated ideas) and code 22 (end of ideation) correlate on some states ($q=2$ or 7) since talking about unrelated ideas may signify that one is situating at the state of running out of good ideas for the brainstorming topic. Another observation is that ideas were *not* clustered and categorized in only one way as presumed by the Markov chain model (Brown et al.,




Figure 2: Emission probabilities of the 8-state HMM

1998). The same idea may be accessed through different mental images with varied emission probabilities. For example, code 10, the idea that “building houses in steep slopes” could be detrimental, was associated with other ideas in different mental images with different semantic relations. In mental image 4 ($q=4$), code 10 was associated with code 5 (“intense rainfall as a factor for DFH”) probably through the inference that intense rainfall on a excavated slope is more likely to cause DFH. While in mental image 6 ($q=6$), code 10 appeared to associate with code 2 (“fragile geological condition”) and code 3 (“deforestation”), and the semantic relation that connect them were different from the relation between code 10 and code 5.

The observation is consistent with Nijstad & Stroebe’s SIAM model that mental images may have fuzzy boundaries and overlap. Our computational model provides support of this view.

Analysis 2: Decoding Idea Generation Sequences

HMM as a computational theory for idea generation makes it possible to decode observed idea generation sequences and examine how mental images transit during idea generation. This allows us to explore what is the relation between the activation/transition of mental images and idea generation performance. There have been two competing views. Some argues that frequent activations of new mental images may lead to worse performance because one may have not generated all possible ideas out of the old images. Also, triggering images from long-term memory may be cognitively more effortful. At the opposite, another view argues that frequent activations of different mental images may be more beneficial because brainstormers may avoid fixation and think more diversely.




Figure 3: Transitional probabilities of the 8 state HMM (showing only arcs with $p > .03$ for clarity)

To explore the plausibility of the two views in the context of the DFH task, we used the Viterbi algorithm to compute the most probable paths of hidden state transitions for idea generation sequence. We then counted the number of between-state transitions for each sequence and correlated it against measures of task performance. For example, if one student's path of mental image transitions across five time slices is [1 1 1 3 5], the number of between-state transition is identified to be 3, which includes the (start)-1 transition, 1-3 transition, and 3-5 transition. To account for the fact that shorter idea generation sequences are inherently more likely to transit less frequently and produce fewer unique ideas, we only included long idea generation sequences (number of total utterances ≥ 8) in the analysis.

Table 2 shows the results of the correlation analysis. There were significant correlations between the number of between-state transitions and the number of unique ideas as well as ideation efficiency (number of unique ideas divided by number of total utterances). In the context of the DFH task, frequent activations and transitions of mental images appear to be an indicator of good ideation performance.

Table 2: Correlations (Pearson's r) between mental state transitions and measures of ideation performance

N=111	A. Num. unique ideas	B. Num. total utterances	C. Ideation efficiency (A/B)
Number of between-state transitions	.648*	-.095	.691*

* $p < .05$

We further compare two students with same number of total utterances but differential performance in idea generation. Figure 3 shows transitional probabilities of the 8-state HMM that we utilize to interpret the two cases. Table 3 shows the decoded paths for the two students. One salient feature of the low performing student (S136), was to transit to mental image q2, and then kept generating unrelated ideas until the end (see Figure 2 for emission probabilities). Another observation was that both students went through mental images q1 and q5 at the very beginning, and then following different paths (either q2 or q4) that seemed to be influential to the final performance. One implication may be the importance of retaining the focus on the task during idea generation. Exploration of off-task ideas (e.g., transiting to q2) might just lead to continuous distraction and end in low number of useful ideas.

Table 3: Profiles of two students

ID	#Total utterances	#Unique ideas	Decoded path of mental image transitions
S208	12	9	[1 5 5 5 4 3 7 7 7 7 7]
S136	12	4	[1 5 5 2 2 2 2 2 2 2 2]

Modeling Group Idea Generation

Dataset and Preprocessing

The dataset consists of ideas on the same DFH task from 34 Taiwanese high school students. The 34 students worked in groups of two, resulting in 17 dyads. Each dyad communicated and generated ideas over Instant Messaging (IM). Ideas were manually coded by using the same coding scheme with 22 codes that we mentioned above. We also assumed that group members took strict turns to contribute. By considering that some members of a dyad may not always use their turns to contribute, and thus the other member of the same dyad may dominate over several consecutive time slices, a new code "empty" was used to denote the situation that one did not contribute anything in her/his turn. Also since off-task social talks were present in these brainstorming episodes, the code "na" was added to mark up that one did not attempt to generate an idea in a specific turn. Among the 17 dyads, 7 were collected and coded in the study presented in Wang et al. (2007). The inter-coder agreement of the coding results reached $Kappa = .84$. The rest 10 dyads were collected in an unpublished follow-up study. The inter-coder agreement of data coding over the 10 dyads achieved $Kappa = .73$.

Note that the size of the dataset is relatively small with respect to the number of free parameters associated with IO-HMMs. For example, if there are 24 possible codes for inputs and outputs, and the number of hidden states is set to 3, the simplest IO-HMM model would still have 284 free parameters that need to be estimated. To control for complexity and increase generality, we reduced the number of possible codes by grouping codes into four general




Figure 4: Modeling dyadic idea generation with direct social influences (with arc between X and Y)




Figure 5: Modeling dyadic idea generation with only influence of cognitive stimulation and interference

categories: valid ideas, unrelated ideas, skipping-the-turn (“empty”), and chitchat (“na”).

Analysis 3: Are Cognitive Explanations Sufficient?

As discussed at the beginning of this paper, one important research topic in the literature of idea generation is searching for explanations of the process loss phenomenon, in which people performed worse when generating ideas interactively in groups than working alone. Early work focused primarily on using social or interpersonal factors (e.g., production blocking, evaluation apprehension) to explain process losses (Diehl & Stroebe, 1987). Recent work focused more on cognitive explanations. Overhearing partners’ ideas could be stimulating if partners’ ideas successfully serve as priming cues for approaching ideas that one has difficulty to access on her own. However, partners’ ideas might also be inhibiting, if they are too similar to one’s own ideas and constrain one’s exploration of unvisited portions of the potential ideation space (Dugosh, Paulus, Rolan, & Yang, 2000).

The relation and interaction between social and cognitive factors remain understated at the current point. It appears that some social factors may be further unrolled as a series of cognitive operations. For example, Nijstad & Strobe (2006) proposed that the effect of production blocking could be explained by disruptions of image activation (i.e., cognitive interference) and idea forgetting when waiting for turns to contribute their ideas. This raises the question that with cognitive explanations at hand, do we still need social




Figure 6: Likelihood-based cross-validation for selecting IO-HMMs

explanations? Do social factors further explain variance that cannot be captured by cognitive explanations?

We trained two types of IO-HMMs with different modeling assumptions to explore this question. In the first type of IO-HMM, the direct-social model, as shown in Figure 4, we assumed that at each time slice one’s idea generation is affected by one’s own mental image and the idea that her partner just verbalized. Note that partner ideas may cause influence directly without the involvement of mental image activation. This type of model therefore represents the scenario that social factors have additional capacity of explanation. In the second type of IO-HMM, the cognitive-only model, as shown in Figure 5, the link between X (partner idea) and Y (self idea) was removed. However, the partner’s ideas may affect the activation and transition of mental images. And idea generation is still governed by mental images. In this way, the influence of partner ideas arrives indirectly. This type of model refers to the situation that social influences are fully mediated by cognitive processes.

We trained n -state IO-HMMs, where $2 \leq n \leq 10$, by using the implementation of EM algorithm included in the Bayesian network toolbox for Matlab (Murphy, 2001). Note that since students of the same dyad produced two interdependent idea sequences, we used the technique of leave-two-out CV to avoid over-fitting in model selection. In each fold of validation, two idea sequences from the same dyad were held out as the testing data, and the rest of 32 sequences were used in training. Therefore, a total of 17 training-testing folds were

executed for each model configuration varying in its structure and the number of hidden states.

Figure 6 shows the goodness-of-fit scores (summations of log-likelihoods derived through CV) associated with each model configuration. We identified the best direct-social model having four hidden states and the best cognitive-only model having seven hidden states. We further compared the two selected models by matching their performance in each fold of the CV procedure. A paired Wilcoxon rank signed test showed that the 7-state cognitive-only model fitted the data better than the 4-state direct-social model ($N=17$, $W_+=49.5$, $p<.05$). Direct social links did not increase the model's modeling capacity. Cognitive explanations appeared to be simpler and sufficient in this domain.

Discussion and Future Work

In this paper, we developed cognitive models for individual idea generation using HMM and group idea generation using IO-HMM. We demonstrated the utility of the models by examining several open questions in the literature of idea generation. It is made possible to examine the semantic relations between ideas emitted by the same image and to decode observed sequences into transitions of mental states. We also took an initial look at the relation between social and cognitive explanations for interactive idea generation via this modeling approach.

We noted the limitation of using only the DFH task (a test-like task that may leave little space for creativity) and the small sample size. The current approach remains at its beginning and needs to be further evaluated in other contexts.

Acknowledgements

The author would like to acknowledge Susan Fussell, Jie Yang, Ziv Bar-Joseph, Yi-Fen Huang, Chao-Lin Liu, Chun-Yen Chang, and I-Pei Tung for their helpful inputs and comments. The research, in cooperation with National Taiwan Normal University, was supported in part by the National Science Council (NSC), Republic of China grant number NSC 94-2511-S-003-028.

References

Bengio, Y., & Frasconi, P. (1996). Input-Output HMMs for sequence processing. *IEEE Transactions on Neural Networks*, 7(5), 1231-1249.

Brown, V., Tumeo, M., Larey, T. S., & Paulus, P. B. (1998). Modeling cognitive interactions during group brainstorming. *Small Group Research*, 29(4), 495-526.

Diehl, M., & Stroebe, W. (1987). Productivity loss in brainstorming groups: toward the solution of a riddle. *Journal of Personality and Social Psychology*, 53(3), 497-509.

Dugosh, K. L., Paulus, P. B., Roland, E. J., & Yang, H. (2000). Cognitive stimulation in brainstorming. *Journal of Personality and Social Psychology*, 79(5), 722-735.

Fan, X. Chen, P-C., & Yen, J. (2007). Learning cognitive load models for developing team shared mental models. *Proceedings of 8th International Conference on Cognitive Modeling*, 145-150.

Feng, G. (2006). Eye movements as time-series random variables: A stochastic model of eye movement control in reading. *Cognitive Systems Research*, 7, 70-95.

Murphy, K. (1998). Hidden Markov Model Toolbox for Matlab. Retrieved, 2007, from the World Wide Web: <http://www.cs.ubc.ca/~murphyk/Software/HMM/hmm.html>.

Murphy, K. (2001). Bayes Net Toolbox for Matlab. Retrieved, 2007, from the World Wide Web: <http://www.cs.ubc.ca/~murphyk/Software/BNT/bnt.html>.

Nijstad, B.A., & Stroebe, W. (2006). How the group affect the mind: a cognitive model of idea generation in groups. *Personality and Social Psychology Bulletin*, 10(3), 186-213.

Raaijmakers, J.G.W., & Shiffrin, R.M. (1981). Search of associative memory. *Psychological Review*, 88(2), 93-134.

Rabiner, L. R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2), 257-286.

Wang, H-C., Rosé, C.P., Cue, Y., Chang, C-Y. Huang, C-C., & Li, T-Y. (2007). Thinking hard together: the long and short of collaborative idea generation in scientific inquiry. *Proceedings of Computer Supported Collaborative Learning 2007 (CSCL 2007)*.

Wang, H-C., Rosé, C.P. (2007). A process analysis of idea generation and failure. *Proceedings of the 29th Cognitive Science Society Annual Meeting (CogSci 2007)*.

Wang, H-C., Chang, C-Y., & Li, T-Y. (in press). Assessing creative problem solving with automated text grading. *Computers & Education*.