Determining What Questions To Ask, with the Help of Spectral Graph Theory

Abe Kazemzadeh[1], *Sungbok Lee*[2,3], *Panayiotis Georgiou*[2], *Shrikanth Narayanan*[1,2,3]

[1]Departmentof Computer Science
[2] Department of Electrical Engineering
[3] Department of Linguistics
University of Southern California, Los Angeles, CA
(kazemzad@,sungbokl@,panos@sipi.,shri@sipi.)usc.edu

Abstract

This paper considers questions and the objects being asked about to be a graph and formulates the *knowledge goal* of a question-asking agent in terms of connecting this graph. The game of twenty questions can be thought of as a testbed of such a question-asking agent's knowledge. If the agent's knowledge of the domain were completely specified, the goal of question-asking would be to find the answer as quickly as possible and could follow a decision tree approach to narrow down the candidate answers. However, if the agent's knowledge is incomplete, it must have a secondary goal for the questions it plans: to complete its knowledge. We claim that this secondary goal of a question asking agent can be formulated in terms of spectral graph theory. In particular, disconnected portions of the graph must be connected in a principled way. We show how the eigenvalues of a graph Laplacian of the the question-object adjacency graph can identify whether a set of knowledge contains disconnected components and the zero elements of the powers of the question-object adjacency graph provide a way to identify these questions. We illustrate the approach using an emotion description task.

Index Terms: dialog agents, question-asking, graph theory, emotions

1. Introduction

Consider a dialog agent that must ask a series of questions in order to identify some unknown object. Whether it be a real, physical object such as a seat on a flight, or a virtual object, such as an emotion or a disease, this object will have a set of attributes that can be described and asked about using natural language. If the agent knows the complete mapping of objects to attributes, it will be able to identify the unknown object by asking a series of questions in the form of a decision tree, where each successive question aims to evenly partition the space, i.e., to "divide and conquer" the space of attributes as quickly as possible. These type of questions satisfy what we call a *task goal*. However, if the agent is acting in a state of incomplete knowledge, it may ask extraneous questions not to identify a particular object, but rather to satisfy a *knowledge goal* [1, 2]. For example, assume that the agent has asked a series of questions that uniquely identify an object, but has incomplete knowledge about some attributes of this object. This agent would continue to ask questions to satisfy a knowledge goal even though the task goal has been completed. One could imagine a conversational diagnostic agent that knows attributes of the flu and attributes of diabetes, but that the attributes do not completely overlap. For example it might know that fevers are associated with the flu, but the association of fever with diabetes is unspecified. Even if the agent has asked enough questions to correctly diagnose diabetes, it may ask another question about fever in order to increase the coverage of its knowledge. It is this type of question-asking behavior–questions motivated by an agent's knowledge goals–that we turn our attention to in this paper.

This paper uses the game of Twenty Questions as a domain to study the knowledge goals of a dialog agent. Although this game is familiar to many and we have observed evidence that it is played in many cultures, we briefly explain the game to establish our terminology. The one player, the *answerer*, picks a word unbeknownst to the the other player, the *questioner*, who attempts to guess the word with twenty or fewer questions. In human-human versions of this game, the two players switch roles at the end of each match. In the formulation we present in this paper, however, we consider the case where the questioner is a computer agent instead of a human. We could also think of a computer agent for the answerer role, but we deem this to be the harder role to automate and look forward to tackling such a problem in future research. The data we analyze is from a specific version of the game of twenty questions that we devised as an experiment to study natural language descriptions of emotions [3]. The basic idea of the Twenty Questions game is still the same except that the answerers must pick words that denote emotions. Despite the limitation of the game to emotion words, the players were not otherwise limited to a fixed emotion vocabulary, in distinction to other emotion research [4, 5, 6] where fixed emotion vocabularies are typically used. This fact, combined with the vague, ambiguous nature of emotion terms led to a difficult task even for the human players. We found that players required on average 12 questions to correctly guess the unknown emotions, when failures to guess correctly are averaged in as twenty questions. We also feel that the limitation of only twenty questions is immaterial from a theoretical and experimental point of view, but practically necessary to prevent undue frustration and expense of time by the players. From other perspectives, the fact that we chose emotions as a domain for questioning is a relevant topic of research which is treated in [3], however in this paper we look at the problem more abstractly. If one considers emotions to be subjectively defined physiological states, such questions could be applied to verbal examinations by doctors, for example. Furthermore, if we consider objects to be general things that can be represented by a vector of attributes, then this model can be applied to a wide range of objects. Emotions, furthermore, have the characteristics of *theoretical objects* [7], i.e., objects whose existence is predicated using natural language that refers to categories of objects rather than actual physical objects. These categories can

The authors would like to thank the participants of our study as well as the NSF, NIH, and DoD.

be seen as equivalence classes formed by the objects' attributes.

The main claim of this paper is that knowledge goals for a question-asking agent can be formulated in terms of completing connections in a graph structure that connects questions with objects. Questions can be thought of as propositions that are assigned true or false values by the answers. These propositions *satisfy*, in a model-theoretic sense, a model of the objects' attributes. However, naively asking questions to connect every question to every answer is inefficient and we are interested in an agent that asks questions that cannot be inferred from others, and we determine these question/object pairs using a graph-based approach. To make the distinction between a lack of knowledge that cannot be inferred and that which can, we define two terms for this purpose. We call lack of knowledge that cannot be inferred from other knowledge *unconnected knowledge* to highlight the fact that it is represented by an unconnected graph. In the case where the lack of knowledge is just due to missing attributes of some objects, we call *incomplete knowledge* to distinguish the graphical representation of this knowledge from a complete graph, where all vertices are connected to every other vertex. Using this graph-based representation, we answer two particular questions: 1) how can an agent determine if its knowledge is unconnected, and 2) how can the agent ask targeted questions that will "connect" the disconnected sets of propositions in its knowledge.

Other work has looked at dealing with uncertainty in dialog systems regarding automatic speech recognition (ASR) output, which results uncertain knowledge. In [8], the problem of determining additional questions to ask is posed in terms of statistical uncertainty about the ASR output and hence the dialog state. In this paper, we simplify the problem by assuming that the dialog system has only one state, that of asking questions about a single object. We also assume that ASR is reliable since it only needs to process answers to yes/no questions. The range of user input is thus limited, so we can assume that ASR only needs to recognize affirmative answers, negative answers, and uncertain answers. At a meta level, the framework of twenty questions to explore a conceptual space has been applied to dialog act taxonomies [9].

2. Methodology

2.1. Constructing a Graph from Question-Object Pairs

In this section, we describe the notion of a theory from mathematical logic which states that a theory Γ is simply a set of sentences in some language $\mathcal{L}$ that is true of a model M [10, 11]. In the case of, Γ is the set of questions that were answered with "yes" and negations of the questions that were answered with "no" for given objects, $\mathcal{L}$ is the language of propositional logic, and M is a model of the objects. Using this formulation, we describe how to construct a graph that can be used to identify important gaps in the agent's knowledge.

In this view, each question can be represented as a proposition p that can be judged true or false of a given object o. The model of a specific object o is denoted M_o. Assuming now that the agent has just asked question p, we can say that $\models_{M_o} p$ if the user answers "yes" or $\models_{M_o} \neg p$ if the player answers "no". The previous notation is read "p/$\neg p$ is satisfied by M_o", or equivalently "M_o is a model of p/$\neg p$". If a proposition p satisfies the model of object o, M_o, then $p \in \Gamma_o$, where Γ_o is the theory of object o. If we can enumerate a complete set of P propositions p_n indexed by $n = 1 \ldots |P|$, then we can represent Γ_o as a Boolean vector of length $|P|$. For every question p_n asked, the n-th position of Γ_o will be *true* or 1 if the user has answered yes to p_n when the object was o. In this case we can say that p_n is a *theorem* of Γ_o. Similarly, *false* or 0 is assigned to element n of $\Gamma_{\varepsilon,i}$ if p_n received no as an answer when being questioned while o was the object in question. In this case, $\neg p_n$ is a theorem of Γ.

If a theory for a specific object can be seen as a long list of propositions that are true of it, the theory of a set of emotions can be seen as an matrix Γ where the rows are indexed by the objects and columns are indexed by the questions/propositions. If the theory Γ contains objects o_m for $1 \leq m \leq M$ and propositions p_n for $1 \leq n \leq N$, then Γ will be an $N \times M$ matrix. Ordinarily, Boolean algebra would dictate that this matrix would consist of ones and zeros. Such a representation has been explored under the aegis of *formal concept analysis* [12]. However, we need the matrix to be sparse to represent the fact that not all of the combinations of questions and emotions have been encountered due to incomplete knowledge. To this end, we propose that the matrix be a $(1, 0, -1)$-matrix, or a *signed matrix/graph*, where 1 indicates that the proposition of column-m is true for the emotion of row-n, -1 indicates that it is false, and 0 indicates that it has not been seen or that a contradiction has been encountered. To make the matrix Γ a square, symmetric adjacency matrix, we define the adjacency matrix of Γ, $A = A(\Gamma)$ to be an $M + N \times M + N$ matrix as follows:

$$A(\Gamma) = \begin{bmatrix} zeros(M) & \Gamma^T \\ \Gamma & zeros(N) \end{bmatrix}$$

This can be seen as saying that questions and the objects they are asked of are both nodes in a bipartite graph. This graph connects questions to objects, and *vice versa*, but does not connect questions with questions nor objects with objects.

The absolute value $|A|$ of A describes whether questions have been asked of objects, regardless of whether the answer was yes or no. It is this graph $|A|$ that gives us information about the connectivity of an agent's knowledge.

2.2. Identifying Unconnected Knowledge

Converting the theory Γ to the graph A, as described above, allows us to use methods from collaborative filtering, social network analysis, and spectral graph theory [13, 14]. In this paper, we use the number of zero eigenvalues of the Laplacian of the graph A to determine the number of connected components of the graph. This can be seen as a measure of the sparsity of our data and can be used to identify the questions that must be asked of certain emotions in order to connect the graph components. The Laplacian L of a signed graph is calculated by subtracting the absolute adjacency matrix $|A|$ from the diagonal absolute degree matrix $\bar{D}_{ii} = \sum_j |A_{ij}|$:

$$L = \bar{D} - |A|$$

From the matrix L we can tell the number of connected components of A by counting the number of zero eigenvalues. Thus, if there are three eigenvalues that equal zero, the graph is composed of three separate connected components. A graph Laplacian with one zero eigenvalue is a single connected graph.

2.3. Determining Which Questions to Ask

If, from Section 2.2 an agent has identified that it has unconnected knowledge, how can it then plan questions to address the knowledge goal of connecting the components of A? To answer this, we must define the notion of a *walk* on a graph. A walk of length l on graph A that joins vertices v_i and v_j is a sequence of vertices $u_0 \dots u_l$ of A such that $v_i = u_0$, $v_j = u_l$, and u_{t-1} and u_t are adjacent for $1 \leq t \leq l$.

According to [15, Lemma 2.5], the number of walks of length l in A that join v_i to v_j is the entry in cell (i, j) of the matrix A^l. Thus, by taking repeated powers of the absolute adjacency matrix $|A|$, we can determine if nodes v_i and v_j are connected by walks of length l. Since the graph is bipartite, the walks from question nodes to other question nodes or from object nodes to other object nodes will always be even length, and conversely, walks between question and object nodes will be odd length. This behavior is undesirable because we wish to preserve connectedness properties across repeated powers of A^l. To remedy this undesirable behavior we can augment the adjacency matrix $|A|$ by adding the identity matrix I to it. At this point, we can say that vertices v_i and v_j are connected by a walk of length l *or less* if the entry (i, j) of $(A + I)^l$ is non-zero. The proof of this, by contraction, that if we imagine that vertices v_i and v_j are connected by some walk of length $k < l$, but not of length l, then there must not be self-loop from v_j to itself after the walk of length k. However, since we added the identity I matrix to A we know that there are in fact self-loops on all of the vertices.

The preceding fact allows us to state an alternative test for connectedness and also allows us to identify the question-object pairs that need to be asked to complete the agent's knowledge. This test can be stated as follows: the graph A is connected if and only if

$$(|A| + I)^{M+N-1}$$

has no zero entries. This is because the length of a walk with distinct steps, a *path*, is at most one less than the number of vertices in the graph, i.e., $M + N - 1$, which would be the case if the graph were a linked list. The question-object pairs that correspond to zero entries in this matrix are precisely the set of candidate questions that need to be asked to connect the agent's knowledge.

The reader may wonder what is the purpose of using the graph Laplacian method to determine whether the graph is connected when this can be accomplished using the method of taking powers of $(A + I)$. While it is true that the latter method can accomplish the same objective of determining whether the graph is complete, the Laplacian eigenvalue method tell *how many* connected components there are. This information is useful because it can tell us how many questions need to be asked. For example, if zero is an eigenvalue of the graph Laplacian, as described in Section 2.2, and this eigenvalue has multiplicity m (i.e., there are m eigenvalues equal to zero), the minimum number of questions that need to be asked is the number of edges to create a spanning tree on m nodes. In this paper, we do not consider any type of weighting on the added edges, so any minimal set of questions that connect the disconnected components of the graph are satisfactory for the purposes of this paper. This minimal set of edges is simply any tree that connects the m disconnected components. Thus $m - 1$ questions must be asked since there are $m - 1$ edges in a tree of m nodes. One could imagine additional constraints that further identify a "best" set of question-object pairs, for example, the set of questions whose added edges minimizes the diameter of the resulting graph.

3. Data

We collected training data for a question-asking agent using a wizard of Oz experiment where humans played both the questioner and answerer roles for the emotion twenty questions game, as described in Section 1 and in more detail in [3]. We collected a total of 26 matches from 13 players. Since each match has two players, this averaged 4 matches per player and ranged from 2 to 12 matches. In the data, a total of 23 unique objects (emotions) were chosen, i.e., only three objects were observed more than once. Table 1 describes how we processed the questions from raw text (of which there was a total of 313 unique questions asked) to a logical representation, which resulted in 222 unique question nodes

Table 1: Data processing

Preprocessing step	Number of questions
0. Raw text	313
1. Text normalization	297
2. Logical representation	222

Since surface forms of the questions vary widely and because at the current stage we have not developed natural language processing techniques to extract the underlying semantics of the questions, we used manual preprocessing to normalize the questions to a logical form that is invariant to the wording. This logical form converted the surface forms to a pseudo-code language with a controlled vocabulary. This standardization involved converting the emotion names to nouns, if possible, standardizing attributes of emotions and the relations of emotions to situations and events. After the standardization, there were a total of 222 question types.

In a basic conversion of our data to a graph, there are a total of $23 + 222 = 245$ objects, our data results in an adjacency matrix A of size 245×245. We will call this the *basic graph*. However, the object identity questions (e.g., "Is it embarrassment?" for our domain of emotion guessing) identify additional objects. These objects have not been picked as an object but it have been referred to in a question. So although we have not seen the object "embarrassment" in our data as the emotion picked by the answerer, it was referred to in one of the questioner's questions. To account for this, we augmented our graph with new vertices for these objects that were heard of but not observed. This simply involved adding more object vertices to the graph and connecting them to the object identity question that referred to them. We call this the *derived graph*. In this case, there were 99 objects, which resulted in an adjacency matrix of 321×321.

4. Results

The Laplacian eigenvalue analysis of the basic object-question graph showed us that there were 35 zero valued eigenvalues, and hence 35 separate subgraphs of our 245×245 adjacency matrix A. Although this shows a high degree of sparsity in the graph, analyzing the repeated powers of $A + I$ showed that all of the disconnected components were all single question vertices that represented infrequently asked questions that were answered without a clear yes or no. Since these were not connected to any of the emotions, asking them of any emotion could serve to connect them to the main part of the graph. However, since these were already asked at least once without a clear yes or no answer, this might actually suggest questions *not* to ask. There were no unconnected object nodes due to four highly used questions that connected all the object nodes. Thus, the power series analysis of $A + I$ showed that there were rows and columns of zeros in elements (i, j) and (j, i) for each question i that was unconnected and the set of vertices corresponding to objects was connected.

The Laplacian eigenvalue analysis of the 321×321 derived graph showed use that this expanded graph actually had the same number of disconnected components, although these components were larger because the unconnected fragments included question-object pairs instead of single questions vertices.

The fact that the number of unconnected components did not increase from the case of the basic graph to the derived graph is an important result. Even though the number of objects more than tripled (from 23 to 99) and no additional questions were asked, the derived graph has similar connectivity characteristics.

5. Discussion

In studying the spectra of the question-object graphs, we were expecting to find large disconnected subgraphs that could be connected with strategically chosen questions. What we actually found was that, in the basic graph, the disconnected subgraphs were trivially single question vertices, which represent knowledge that is only unconnected in the question vertices. The main body of the graph was connected due to several frequently used questions. Therefore, the answer to the question that we set out to answer, "what questions should an agent ask when dealing with unconnected knowledge?", is that the agent could ask any of these questions of the any objects since in the basic graph the objects are all connected. One possibility that must be considered though is that these questions may be disconnected for a reason: they could be irrelevant. However, the data we based our graphs on is from human-human interaction data, so we must assume that these disconnected questions were relevant *for the players who asked them*. We discuss player-specific theories at greater length in [3], the conclusion of which is that it is useful to model individual players who may differ in knowledge and question-asking strategies. Thus, if an agent is designed to play Twenty Questions as a human would, then it makes sense to ask these questions, while for best task performance it it would be best for an agent to simply prune away these disconnected nodes. For modeling human performance these questions capture the fact that humans are not optimal question askers with complete knowledge, at least for our task of asking questions about emotions.

In the derived graph, which had additional vertices added for unseen objects that were referred to in questions, the situation was different. The derived graph had disconnected object vertices as well as question vertices. If the agent is free to pick any object-question pair to generate a question, then whether the unconnected vertices are questions or objects is not important. However, in the Twenty Questions game, the other player chooses the object, so the agent will have less opportunities to ask questions about specific objects for which its knowledge is unconnected. However, if the object is connected to other questions that are also disconnected from the main body of knowledge, then the agent can ask these questions about any other object that is connected with the main body of knowledge because this will result in the unconnected object being connected via the question.

6. Conclusion

This paper presented a way for a question-asking agent to deal with incompleteness in its knowledge. We found that the eigenvalues of the graph Laplacian and the power series of the adjacency matrix give information that an agent can use to determine which questions to pick in order to complete it's knowledge. In particular we found that the knowledge that was unconnected in the basic object-question adjacency matrix was knowledge about individual questions. When we included data about unobserved objects, those that had been asked about but not seen, we found that the object nodes were also disconnected, but that the number of disconnected components remained the same.

In this work, we considered only questions that were observed in our data. However, questions are often of a type that can be generalized. For example, "Is the emotion similar to happy?", can be generalized to "Is the emotion similar to X?" where X could be any word for an emotion. Expanding questions of these types would lead to even more extensive incomplete knowledge, but could lead to better abilities to make inference between the questions, an issue that we will examine in future work.

7. References

[1] A. Ram and D. B. Leake, *Goal-driven learning*. MIT Press, 1995.

[2] A. Gordon, *Strategy Representation: An Analysis of Planning Knowledge*. Lawrence Earlbaum, 2004.

[3] A. Kazemzadeh, P. G. Georgiou, S. Lee, and S. Narayanan, "Emotion twenty questions: Toward a crowd-sourced theory of emotions," 2011. Under review.

[4] A. Kazemzadeh, S. Lee, and S. Narayanan, "An interval type-2 fuzzy logic system to translate between emotion-related vocabularies," in *Proceedings of Interspeech*, (Brisbane, Australia), September 2008.

[5] A. Kazemzadeh, "Using interval type-2 fuzzy logic to translate emotion words from spanish to english," in *IEEE World Conference on Computational Intelligence (WCCI) FUZZ-IEEE Workshop*, 2010.

[6] A. Kazemzadeh, S. Lee, and S. Narayanan, "An interval type-2 fuzzy logic model for the meaning of words in an emotional vocabulary," 2011. Under review.

[7] T. Forster, *Reasoning About Theoretical Entities*. World Scientific, 2003.

[8] J. D. Williams, P. Poupart, and S. Young, "Partially observable markov decision processes with continuous observations for dialogue management," in *Proceedings of SIGDIAL Workshop on Discourse and Dialogue*, (Lisbon, Portugal), October 2005.

[9] D. Traum, "20 questions on dialogue act taxonomies," *Journal of Semantics*, vol. 17, no. 1, pp. 7–30, 2000.

[10] T. Forster, *Logic, Induction, and Sets*. Cambridge University Press, 2003.

[11] H. B. Enderton, *A Mathematical Introduction to Logic*. Academic Press, 2nd ed., 2001.

[12] B. Ganter and G. S. R. Wille, eds., *Formal Concept Analysis: foundation and applications*. Berlin: Springer, 2005.

[13] J. Kunegis, A. Lommatzsch, and C. Bauckhage, "The slashdot zoo: Mining a social network with negative costs," in *World Wide Web Conference (WWW 2009)*, (Madrid), pp. 741–750, April 2009.

[14] Y. P. Hou, "Bounds for the least laplacian eigenvalue of a signed graph," *Acta Mathematica Sinica*, vol. 21, no. 4, pp. 955–960, 2005.

[15] N. Biggs, *Algebraic Graph Theory*. Cambridge University Press, 1974.