Models of the Behavior of People Searching the Internet: A Petri Net Approach

Paul B. Kantor

APLab, SCILS, Rutgers University
4 Huntington St., New Brunswick, NJ   08903
kantor@scils.rutgers.edu

Ragnar Nordlie
Oslo College, Oslo, Norway
ragnar.nordlie@jbi.hioslo.no

Abstract

Previous models of searching behavior have taken as their foundation the Markov model of random processes. In this model, the next action that a user takes is determined by a probabilistic rule which is conditioned by the most recent experiences of the user. This model, which has achieved very limited success in describing real data, is at odds with the evidence of introspection in a crucial way. Introspection reveals that when we search we are, more or less, in a state of expectancy, which can be satisfied in a number of ways. In addition, the state can be modified by the accumulated evidence of our searches. The Markov model approach can not readily accommodate such persistence of intention and behavior. The Petri Net model, which has been developed to analyze the interdependencies among events in a communications network, can be adapted to this situation. In this adaptation, the so-called "transitions" of the Petri Net occur only when their necessary pre-conditions have been met. We are able to show that various key abstractions of information finding, such as "document relevance", "a desired number of relevant documents", "discouragement", "exhaustion" and "satisfaction" can all be modeled using the Petri Net framework. Further, we show that this model leads naturally to a new approach to the collection of user data, and to the analysis of transaction logs, by providing a far richer description of the user's present state, without inducing a combinatorial explosion.

INTRODUCTION

It was proposed some twenty years ago (Penniman, 1975) that models of the behavior of users of search systems could be helpful in developing systems that are more effective in serving the user. The overall argument goes like this: if we have an effective model of the behavior of the user of the system, then when we observe deviations from this effective model we can recognize that something strange is happening to the user and invoke some sort of adaptive or corrective mechanism. In the early days this was imagined to be some kind of help demon that would pop up and provide a suggestion. The idea remains an attractive one and forms the basis of new research proposals every year. The idea, however, stands or falls on the existence of an effective model describing the behavior of the user of a search system. To be effective the model must be fairly accurate. Since human beings are not automata, we can not expect that the model will be perfectly accurate. However, it needs to make some kind of prediction about what the user will do next, given information about what the user has done already. In fact, the prediction it makes must be sufficiently sharp that we can also algorithmically decide whether the behavior of the user has deviated from the prediction. This means that we must generate a mathematical approximately predictive model of the user's behavior.

This technical requirement means that some models which are conceptually very interesting, such as a model which characterizes a search in terms of cycles of behavior, recurrent through a special search behavior or state thought of as a fresh start, are not satisfactory for this purpose. A typical model of this form regards repetitions as a non-event. In other words, it doesn't actually make a prediction as to whether

the next thing that a user will do is to continue in the present part of the cycle, or to move to another part of the cycle, or perhaps start a new cycle completely.

In the search for mathematical models that can make predictions, the field has been dominated by the initial work of Penniman (1975,1982) and Tolle (1983)who postulated a so-called "first order Markov model". Markov models, although they can be quite complex, can be summarized in a very simple way. When we describe a user by such a Markov model we say that the user (or the user's behaviors) can be classified into a relatively small number of well-defined groups called "states". The Markov model says that when a user is in one of those states, the users next action will be to move to one of the other states, according to a fixed random law. The notion of a fixed random law means that for any particular state there is a definite set of probabilities governing the chance that any particular state will be the next one. For example, if there are three states in the system (new start, save results, continue present search), which we will represent in shorthand as ($N,S,C,$), a typical Markov model might say that if a person is in situation "C" the chance of going to each of the three states is given by some probabilities whose sum is 1, such as (20%, 30%,50%). These probabilities mean that if a searcher is continuing (that is, has just continued at his most recent step,) then the chance that he will continue at the next step is given by 50%. The chance that he will make a new start is 20%, and the chance that he will save his search results is 30%. Subject to all the usual difficulties of dealing with probabilities, this is a very concrete and specific model. Over a reasonably short term, if a user were generally described by such a model, it would be possible to see that he has deviated from his usual behavior.

Unfortunately, however, these first-order Markov models, fail almost immediately. This can be verified as follows. If any particular behavior is governed by a first-order Markov process, then the chance of observing a string or chain of length l decreases geometrically with the length of the chain. In other words, if there's a 50% chance that a continuation will be followed by a continuation, that means that there's a 25% chance that a continuation is followed by two continuations and there's a 12.5 % chance that is followed by three continuations and so forth. It is easy to take almost any set of search behavior data, isolate the chains of repeated events, and verify that they do not follow a geometric law. In other words, the first-order Markov model fails.

There are three different approaches that can be taken when a model of this type fails. The first which has been explored both by Penniman (1982) and Chapman (1981) is to go to higher order Markov models. In a higher order model, for example a second order model, the description is still in terms of probabilities, but the probability that my next step will be a continuation depends not only on my present state, but also on the immediately preceding state. Models of this type can be built up to be more and more complex, and there has been some effort apply them particularly by Qiu (1993). In a general way the best we can say is that these models are not yet really satisfactory.

There are two other directions in which research on salvaging the Markov idea could be pursued, although they have not yet been. The first is to redefine the set of states. States have been chosen on the basis of some qualitative understanding of the search process, and the choice is not really validated by any deep cognitive understanding of the way people search. So perhaps if we were to replace our present choice of states by others we would see that what looks to us now like a long chain in exactly the same state really represents a series of transitions among other states.

The second way of extending a Markov model is to ask whether there are some states which are "amnesiac". This is exactly what the Markov model would like. When I fall into one of the amnesiac states I *do* forget where I've been and my choice of a next state can be accurately described by a random model. On the other hand, there would be states which are "memoryful" in which a person really does remember what has happened before and so the transition or the state following that state is determined by the deeper history. While there is much to be said for random models, in terms of their clarity and simplicity, and while they have been used with success to describe some kinds of human behavior (voting behavior, purchase behavior, and so forth) we believe that another perspective provides a better foundation for tackling the issue of modeling human searching behavior.

A MEMORYFUL PERSPECTIVE

From the point of view of introspection, which is available to all of us, the idea that we do what we do because we remember the past seems entirely natural. In other words, it would seem to us that in trying to fit Markovian models to our human behavior we're really starting at the wrong end of the spectrum. We ought to start with something that recognizes that we do remember the past. To remember or to know the past does not mean that you have to lack volition or be capricious in your behavior. There is a very precise technical sense in which physical systems remember the past excellently. In other words, while the present state of motion of all the particles in the world is supposed to be able to form the basis for predicting the future motion, it's also true that that present basis is supposed to be able (under a very deep principle called time reversal invariance) to let us retrodict the entire past. In other words, that present state of information contains complete knowledge about the past as well as whatever it can predict about the future.

Guided to some degree by this kind of thinking we can propose what might be called a First Law of Motion relating to searching behavior. Newton's First Law, which describes the motion of inert physical masses, states that a body at rest tends to remain at rest and a body in motion tends to remain in motion. In much the same way we would like to sharpen the idea that a person engaged in searching tends to remain in motion, whatever that may precisely mean, more or less in the same way that he is in motion now.

In the world of physics obviously things don't remain at rest or in uniform motion all the time. In fact all kinds of interesting things happen. The way this is accommodated in physical law is to say that something called "forces" act on particles or masses and cause them to change the way they are moving. In a very simple analogy, we propose that there are various kinds of cognitive and informational forces which act on a human searcher and cause her to change her behavior or motion.

In thinking about people as they are searching, we want to look for a model which captures the notion that there is a kind of momentum in the search process and that the key events are changes that deflect the course of the search process. In the physical analogy these changes are caused by external forces. In the case of humans searching we propose that they be thought of as transitions, caused by or occurring in response to external events.

There is a quite well developed mathematical model for discussing systems whose key events are transitions. This model, called the Petri Net Model (Reisig, 1985), was first developed to study problems of synchronization and timing in computer networks and has since been expanded. The central concept of the Petri Net Model is the transition. A transition can be thought of as a node or point in some kind of a network. Every transition has associated with it a set of abstractions called its "input places" and another set of abstractions called it "output places". Under the basic rules of Petri Nets a transition "fires" if and only if there is at least one marker or token in every one of its input places. In usual operation, when a transition fires it removes exactly one token from each of the input places and distributes one token (tokens are not real things, they are abstractions and the number of tokens distributed may be quite different from the number of tokens removed) to each of the output places.

It is easy to see that this basic framework can accommodate the description of many of the kinds of transitions that may occur during searching. For example, we could have a transition called "quit" which has one input place which is called "find a good document". This model rule means that when a good document is found, a token is placed in the "find a good document" place, the quit transition realizes that its firing conditions are satisfied and it fires. To complete the cycle we could say that there's a place called "quit" and the "quit" transition places a token in it.

In fact an enormous mathematical theory has been developed around these apparently simple concepts, in which the state of the system is a very complex structure represented by a complete listing of how many tokens there are in each of the places of all of the transitions in the network. This complexity is encouraging, because we like to think that mental states of human searchers are very complex, even though we hope to describe them in terms of a relatively simple number of basic transitions.

An old idea in information retrieval is that a person may be looking not for a single good or relevant document but for some specific number of them, say 5. The reader may wish to try the exercise of developing a Petri Net using places and transitions as we've just defined them to describe a person who keeps searching until 5 good documents have been found, and then quits. However, we warn you in advance that you won't be able to do that. It turns out that some extension of the Petri Net concept is required in order to do basic arithmetic computations. The simplest way to make the extension is to add two more concepts. The first concept is multi-token removal. If a transition cannot fire until there are, for example 5 tokens in one of its input places (and perhaps just one in each of the other input places) then we can do the kind of counting that we just described. We would just say that this person has a link of weight 5 between the input place for good documents and the transition representing quitting.

For more complicated calculations, it turns out also to be necessary to introduce a suppressor link. What this means is that a particular input place may have a negative effect on its transition. So, if this place is occupied, the transition cannot occur. This can be used to hold one transition off and prevent it from occurring until a certain other trigger transition has occurred. This is shown in Fig. 7 where transition A will take a token out of the place called "S" and when that token is removed, transition B will be free to fire.

With these added complexities it turns out that the Petri Net Model becomes very rich indeed. In fact, it has been shown mathematically that extended Petri Nets, including these two additional concepts, are able to perform all the computations that can be performed by a Turing Machine; in other words, they are general purpose computers of a very high order.

## MODELING SOME COMMON SEARCH BEHAVIORS BY PETRI NETS

Basics of Petri Nets

Figure 1. A basic transition
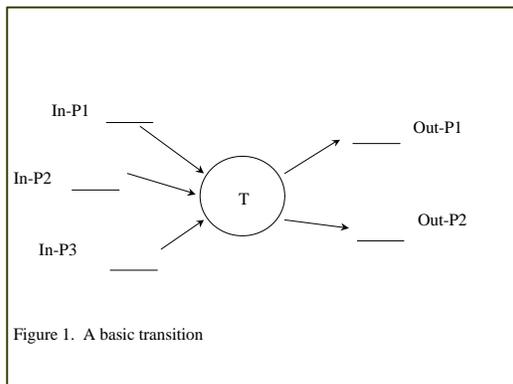
The central notion of Petri Nets is the dynamic concept of a transition, represented by a circle. (Figure 1.) A transition (T) is an abstraction of the notion that some event, represented by T, occurs only when other events, on which it depends, have occurred. Those events are represented by *input places*, shown on the left. When the transition occurs, it establishes that some new set of conditions have been satisfied. These are represented by the *output places*, to the right of the diagram

Figure 2. A transition whose conditions for fining are met.
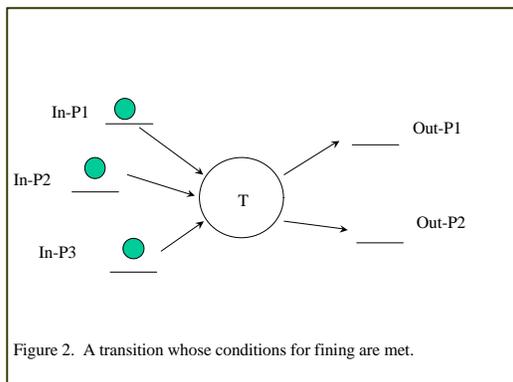
As the conditions represented by the various places are met, this is indicated by tokens (represented by the shaded circles) in the places. The transition shown in Figure 2 is ready to fire, as all of its input places are filled.

Figure 3. The transition of Fig. 2 has fired.
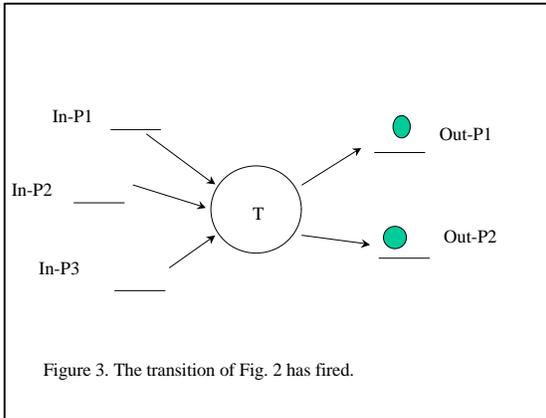
In Figure 3, the transition has fired, and tokens now appear in the output places.
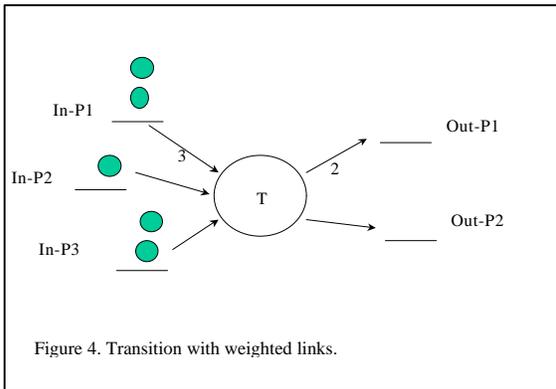

Figure 4. Transition with weighted links.

Figure 4 shows that the links in a diagram can be given weights. For example, a transition with an input place having a link of weight 3 will not fire until there three tokens in the place, and when it does, it will remove all three of them. Note that if an input place with link weight 1 has accumulated 2 tokens, there will be one left. The transition in Figure 4 is ready to fire when one more token arrrives in input place In-P1.

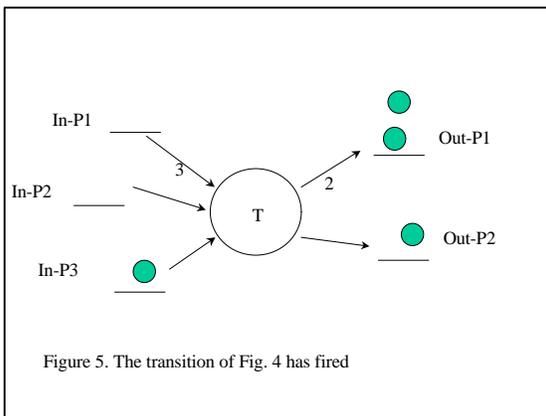Figure 5. The transition of Fig. 4 has fired

Figure 5. After one more token arrived to input P1 of Figure 4, the output link of weight 2 results in the appearance of two tokens in the place Out-P1. Note that one token has been left in place In-P3.
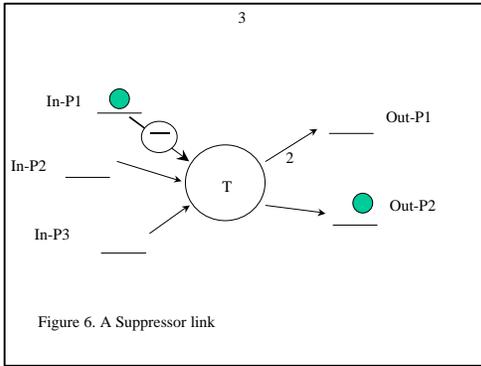
Figure 6. A Suppressor link

Figure 6. An even more powerful concept is the suppressor link, indicated by the minus sign in Figure 6. As long as there is a token in place In-P1, the transition T cannot fire. This powerful tool permits the use of Petri nets as generalized arithmetic and decision making devices, in a way that can represent many kinds of human behavoir.

Representative Simple User Models

Figure 7. The Petri Net for stopping
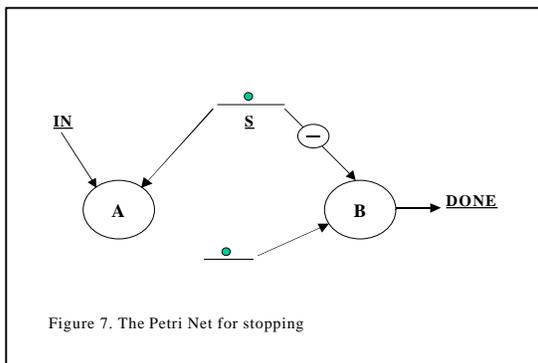
In Figure 7 we show the key elements of a Petri Net representing the decision to stop some activity. There are two transitions. The one on the right, which results in a change, is suppressed by the token in the place labeled "S". But as soon as the place labeled "IN" acquires a token, the transition on the left will fire, removing the suppression token in place S.

Figure 8. A user seeking 5 good documents.
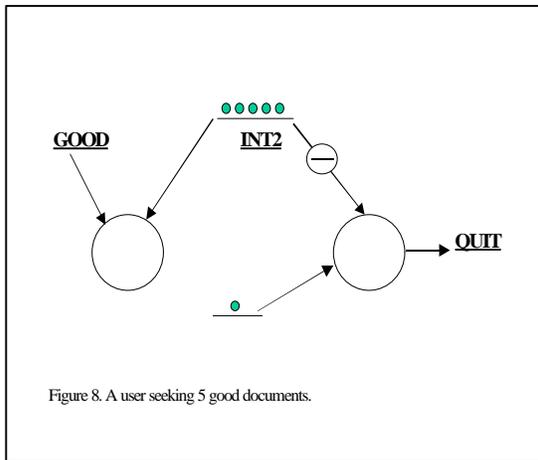
With these basic principles we can build simple networks of transitions to represent possible rules used by human searchers. For example, the network of Figure 8 is a generalized network for a searcher who want to find 5 relevant documents, and will quit when, and only when she finds them all.

Figure 9. A user with bounded optimism
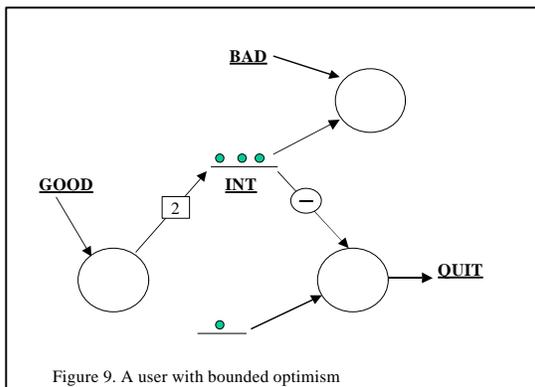
More sophisticated model allow for other circumstances under which user's change their behavior – in other words, respond to external stimuli or forces with a change in search direction. Still considering that inputs are simply relevant or not relevant, we can model some other kinds of behavior. For example, users may base their decision on some rough estimate of how likely they are to find additional useful documents.

Here we show a user who will quit if either of two circumstances occurs: starting off with 3 bad (that is,

not relevant or encouraging) documents, or developing a running average in which more than 1/3 of the documents are not relevant. The QUIT place will be occupied if the first three documents examined are "BAD", as that will exhaust the tokens in the intermediate place. But if a good document is seen, it will add two more tokens to the place, prolonging the time until the user quits.

One possible evolution of the model shown in Fig. 9 is detailed in Table 1.

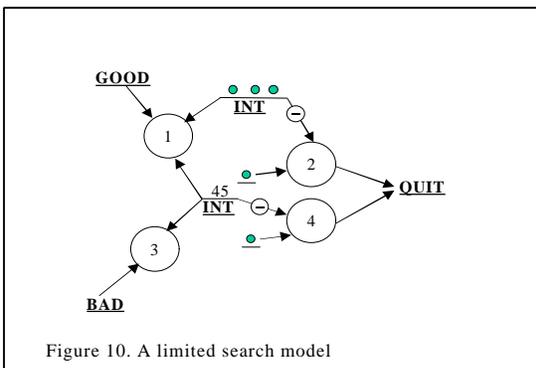| Table 1 A possible evolution of Figure 9.				
Status	BAD	GOOD	INT	QUIT
Initial	0	0	3	0
See a bad doc.	1	0	3	0
	0	0	2	0
See a bad doc.	1	0	2	0
	0	0	1	0
See a good doc.	0	1	1	0
	0	0	3	0
See a bad doc	1	0	3	0
	0	0	2	0
See a bad doc.	1	0	2	0
	0	0	1	0
See a bad doc	1	0	1	0
	0		0	1

Figure 10. A limited search model

Figure 10 models a searcher with a finite appetite for the entire search process. As soon as three good documents have been seen, the upper intermediate place is empty, and transition "2" can fire, leading to "QUIT". In addition, however, each document seen, whether good (relevant) or bad, removes one of the 45 tokens in the lower intermediate place. And when 45 documents have been seen, the lower transition is free to fire, and does so, resulting in a "QUIT".

These basic elements show that the reasonable parameters of "expectancy" can be modeled using Petri nets. As the documents are examined, the user stays in qualitatively the same state (for example, continue) until some crucial accumulation of evidence or influences causes the final transition to a major different state (in this case, QUIT). In other words, the intuitively defined states representing user behavior are collections of the more precisely defined states. A user with no token in QUIT, but with only 10 tokens in the central intermediate state is still "CONTINUING", but examination of the interior state detail shows that he is "running out of patience".

PROBLEMS OF APPLICATION

Up to this point we have presented only a theoretical or conceptual framework for describing the behavior of searchers. We find it philosophically appealing because it shifts emphasis from random jumping from one state to another, as in the older Markov Models, to fact-driven behavior consisting of transitions among various kinds of persistent search activity.

In order to confront this theory with reality we need to be much more specific about the map between these abstractions and behavior of real people as they search. In doing this we have to operationalize some fundamental kinds of transitions. Our thoughts on this are still somewhat speculative. We believe that in the course of searching in a networked environment there are several kinds of transitions that are apparent to the user. One is to give up. A second is to change the way in which you think of the question. A third is to persist with the same question but to use some kind of a search engine and describe your same quest in new terms or key word (reformulation). A fourth is to find a particular item to be useful for navigation. A fifth is to find a particular item to provide some useful information. A sixth is to find a particular item to not provide any useful information.

While these alternatives seem conceptually clear, it is not obvious that we can elicit information from real searchers about which of these events or transitions has just occurred to them. In fact, it is not apparent that searchers can be enabled to recount their search experiences in a language of persistences and transitions. We propose here to embark on an empirical program of investigation which builds on these theoretical ideas and seeks to determine whether Petri Net, transition-based models of searcher behavior can provide an effective characterization of human searching behavior.

Due to very rapid advances in technology, we are planning to conduct research using screen capture and post-search interview techniques which we believe will be minimally disruptive to the searchers mental processes. We are presently developing a tools and instruments for the formalization of this process and they will be described in subsequent reports.

REFERENCES

Chapman, J.L. (1981). A state transition analysis of on-line information-seeking behavior. *Journal of the American Society for Information Science,* 32, 325-333.

Fenichel, C.H. (1981). Online searching: measures that discriminate among users with different types of experiences. *Journal of the American Society for Information Science,* 32, 23-32.

Penniman, W.D. (1975). *Rythms of dialogue in human-computer conversation.* Unpubl. Doctoral diss., Ohio State University, Columbus, OH, 1975.

Penniman, W.D. (1982). Modeling and evaluation of on-line user behavior. *Information interaction: Proceedings of the 45th ASIS information meeting, Columbus, OH, October 17-21, 1982.* White Plains, NY: Knowledge Industry Publications, 1982.

Qiu, L. (1992). Markov models of search state patterns in a hypertext information retrieval system. *Journal of the American Society for Information Science,* 44(7), 413-427.

Reisig, Wolfgang (1985). Petri nets : an introduction. Berlin ; New York : Springer-Verlag,.

Tolle, J.E. (1983). Current utilization of online catalogs: Transaction analysis. *Final report to the Council on Library Resources, 1.* Dublin, OH: OCLC, 1983.