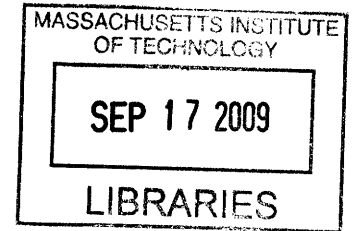


Visualizing the Conversation Pathways of Telephone Linked Care in a Directed Graph

by

Penelope H. Cuevas

B.S. Biology
Massachusetts Institute of Technology, 2002
M.D.
Albert Einstein College of Medicine, 2007



SUBMITTED TO THE DEPARTMENT OF HEALTH SCIENCES AND TECHNOLOGY IN
PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE IN BIOMEDICAL INFORMATICS
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2009

ARCHIVES

©2009 Penelope H. Cuevas. All rights reserved.

The author hereby grants to MIT permission to reproduce
and to distribute publicly paper and electronic
copies of this thesis document in whole or in part
in any medium now known or hereafter created.

Signature of Author: _____

Department of Health Sciences and Technology
May 18, 2009

Certified by: _____

Robert Friedman
Chief, Medical Information Systems Unit, Boston University
Thesis Supervisor

Certified by: _____

Peter Szolovits
Professor of EECS & HST
Thesis Supervisor

Accepted by: _____

Ram Sasisekharan, PhD
Director, Harvard-MIT Division of Health Sciences and Technology
Edward Hood Taplin Professor of Health Sciences & Technology and Biological Engineering

Visualizing the Conversation Pathways of Telephone Linked Care in a Directed Graph

by

Penelope H. Cuevas

B.S. Biology

Massachusetts Institute of Technology, 2002

M.D.

Albert Einstein College of Medicine, 2007

SUBMITTED TO THE DEPARTMENT OF HEALTH SCIENCES AND TECHNOLOGY
ON MAY 18, 2009 IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE
DEGREE OF MASTER OF SCIENCE IN BIOMEDICAL INFORMATICS

ABSTRACT

Telephone linked care (TLC) is a telehealth intervention that has been shown to be effective in a variety of clinical settings. TLC is an interactive computerized telephone system. The system ‘speaks’ to patients, asking them questions and understanding their responses. There is logic built into the calls, so that a patient’s response to a question will dictate the next question that is asked. This serves to personalize the call for each patient, and makes the conversation more realistic. All of the patients’ responses are stored in a database. This database provides much opportunity for analysis because a single phone call contains many responses.

Visualization is an important way of gaining insight into data. Visualization can make data easier to understand and process. Different aspects of data can be encoded in a visualization. The TLC data lends itself to visualization. By viewing each of the questions that the system asks as nodes, and connecting the nodes by the chronological order in which these questions are asked, a tree structure will reveal the conversational paths that are taken in the calls. By combining data from multiple calls and encoding them in this tree structure, new insights can be gained into the TLC data. For example, the frequency with which questions are answered in a particular way can be encoded to reveal the most common pathways through the tree.

This paper describes a visualization application of TLC data which will allow researchers to gain new insights into the TLC conversations and into medical interviews in general.

Thesis Supervisor: Robert Friedman

Title: Chief, Medical Information Systems Unit, Boston University

Thesis Supervisor: Peter Szolovits

Title: Professor of EECS & HST

INTRODUCTION

Patient interviews are a key part of medical care. They are necessary to obtain information about a patient for the purpose of diagnosis and treatment. When patients are interviewed, the interviewer usually has a specific goal in mind. Often the patient comes in with a specific complaint, and the goal of the interviewer is to gather enough additional information to make a diagnosis. Sometimes, interviews are conducted to gain a general sense of the health of the patient for preventive purposes, or the interview might be done to engage the patient in a discussion about issues that are important to their health. Once an interview is complete, and the objective of the interview has been met, the interviewer usually takes away only the information relevant to the initial goal of the interview. However, there is potentially a wealth of information that could be gleaned from the content, flow and structure of the interview if the interview is examined as a whole, and not looked at merely as a means of achieving the goal.

To explore this idea, a specialized patient interview tool called telephone linked care can be used as a good interview model. Telephone-linked care (TLC) is a telehealth intervention which aims to automate the patient-provider interaction and remove some of the disease assessment and management burden from overstretched care providers. TLC is a computerized telephone interview or conversation. Patients can either call into the TLC systems, or the system can call them. The system speaks to the patient via an automated voice, either a recorded voice actor, or through computer-generated text to speech. Patients respond to the system either by pressing buttons on their telephone keypad or by speaking into the telephone where their responses are decoded by speech recognition software. Answers are restricted to a finite answer list (patients are told what their available answer

choices are). Often the questions are yes/no questions, but other answers types are also possible.

TLC systems are programmed to ask questions, and understand patients' responses. Then, a patient's response to a given question will guide future questions to be asked. When the call begins, patients first go through an introductory section where they verify their identity. They then enter the main section of the call where clinical content (in one of many possible disease arenas) is discussed. The clinical content could be focused toward disease assessment and management, behavior modification and counseling, education or combinations of these. The call then ends with a closing section which may summarize the call and establish goals for the future (1).

TLC systems have been evaluated in a variety of disease settings including alcohol use disorders, cancer, nutrition and exercise, and have been shown to have positive outcomes (2, 3, 4, 5).

During the TLC phone calls, all responses given by the patients are time stamped and stored in a database. One can then query the database, and use the caller's responses to reconstruct the conversation that occurred in a given call.

Because of the nature of TLC systems, they offer a valuable opportunity to gain further insight into the flow and structure of the medical interview. This is a relatively unexplored area – since the TLC systems are designed as a telehealth intervention, they are usually just evaluated as a whole (patients who used the system are compared with those who didn't to see if they had different health outcomes) without looking further into the interaction that patients have with the system. However, since all of the responses that the patient gives are stored in a database, a 'snapshot' of the entire interview is available. By

looking at this snapshot for each call and comparing it to those of other calls, new insights into the TLC systems, and the medical interview in general, can be gained. There are many different directions across which the call snapshots can be compared. Patients can be divided into groups based on demographic or clinical variables, and then different groups of patients can be compared to each other. Individual calls can be compared to each other and to combinations of other calls. Calls can be compared across time (for many TLC systems, the same patient will complete multiple calls).

There are limitless ways in which the call ‘snapshots’ could be portrayed or visualized. One way is to visualize a call in a mathematical representation (a graph) that is displayed in a visual format (node-link diagram). A graph is a type of mathematical structure which contains objects (nodes), and connections between the objects (edges). In directed graphs all edges point in a given direction, giving an order to the nodes (Figure 1). This type of visualization is used to represent many types of data, from biological processes to social networks.

The directed graph is a very valuable way to visualize the TLC data. Each question asked during the TLC interview is represented by a node. Edges show the chronological order in which the questions were asked – so that an edge between two questions implies that the second question was asked immediately after the first (Figure 2).

By varying the thickness of the edges, the graph can show how frequently different pathways are taken through the calls (Figure 3). In addition to edge thickness, other visual attributes can increase the degree of information shown. By using color to distinguish different pathways, different groups of patients can be compared (for example, to see if men and women follow different pathways, or older patients and younger patients differ in

their call patterns). Color can also be used to compare the pathways of different calls made by the same caller. Patients can also be divided into groups based on attributes determined in the call – so for example, color can also be used to compare the pathways of patients who answered yes to a given question to those of patients who answered no to that question.

This paper describes an application which displays TLC data in a directed graph. Only the characteristics of the edges in the visualization are used for encoding information. The position of the nodes, and therefore the overall shape of the graph do not change. Thus the basic graph will look the same, allowing for easy comparison between groups.

The visualization application is interactive to encourage users (TLC researchers or others) to explore the data. An interactive visualization tool that is easy and enjoyable to use and aesthetically pleasing rewards users for analyzing the data. The visualization is easy to use, so that even a novice user can gain new and valuable insight into the data.

METHODS

The visualization application was developed using the Processing environment (6).

The first task in creating the visualization was to create a graph showing all of the conversation pathways. As described by Bastert and Matuszewski (7), the optimal graph will have all edges short and pointing downward, have evenly distributed nodes and few edge crossings. In addition to these general principles, for the purpose of the application, the placement of the nodes must be such that they can all fit on one computer screen.

To achieve this, the vertical position of each question node is assigned first: each question node is assigned to a layer. After all nodes have been assigned to a layer, the number of layers is taken into consideration, and each node is assigned a y-coordinate such

that all layers will be evenly distributed in the visualization window. The layering assignment is done by performing a depth first search of the nodes. First, all root nodes are identified (questions that are asked first in the call), then each subsequent question node encountered in the search is assigned to the layer immediately below the previous question. Because a given question could be asked as a part of multiple pathways (there are possible cycles in the graph), nodes that have been previously assigned a layer, are not reassigned.

After being assigned a vertical position, horizontal position assignment is done in a multi-step process. Again, it is done with the goal of equally spacing the nodes while fitting them onto the computer screen. First, helper nodes are added to connect any neighboring nodes that are not in adjacent layers. One helper node is added to each layer between the two nodes. Helper nodes are also added to prevent any edges from overlapping. Then, within each layer, nodes are re-ordered to reduce the number of crossings between edges using the Barycenter heuristic. Then, layer by layer, each node is assigned an x-coordinate. Edges are drawn so that they are as close as possible to a straight line connecting each pair of nodes.

To evaluate the application, it was demonstrated to a focus group of TLC system designers and informaticians. Participants were asked for their qualitative feedback.

SYSTEM DESCRIPTION

Figure 4 shows the application window. There are four main sections of the application window. The large square on the left is the main window of the visualization application. This is where the graph is displayed. Each question node is drawn as a circle. Edges are drawn as lines connecting the circles – thicker lines indicate more frequently

taken pathways. Arrowheads indicate the direction of edges, showing the chronological order in which questions were asked.

On the right side of the application is the control area. The top of the control area contains a smaller version of the overall call graph. Plus and minus signs next to this smaller graph allow users to zoom in and out, respectively. This is useful when the number of questions is large and one particular area of the graph is of interest. It is possible to use this to magnify the graph up to ten times its original size. When the graph is zoomed in, a gray rectangle over the small graph shows which portion of the graph is being viewed in the main window. The user can use the mouse to move this rectangle which will change the portion of the graph shown. Depending on the complexity of the original data (which is determined by the complexity of the particular TLC system), zooming may be unnecessary, as in the very simple graph shown in Figure 4.

The second part of the control area shows users information about the individual questions. This area is blank unless a particular question is selected. As shown in Figure 5, when a user clicks on a question node in the main window, that node changes color, to indicate that it is the current focus. More information about the selected node is shown in the second part of the control panel (the middle section on the right side of the application). The text of the selected question is shown, as well as a bar graph listing the possible responses to that question. The bar graph shows the number of people who gave each response. By clicking on a bar (Figure 6), a user can highlight the pathways – and relative frequency with which the pathways were taken – of callers who answered a question in a particular way. The overall graph remains the same, but all pathways which include the selected response to the selected question are shown in red. This allows the user to see if

callers who answer a particular question in a certain way also answered other questions similarly throughout the course of the call.

The final piece of the control panel allows the user to select individual callers or caller attributes to view. As shown in Figure 7, users can select different caller characteristics (e.g. gender, diagnosis). Pathways of callers with these attributes are then highlighted in red. Figure 8 shows that the pathways of individual callers can also be viewed. The arrow to the right of the ID number allows users to browse through the different callers. For a given caller, one or more individual calls can be viewed.

Figure 9 displays a more complicated graph, and shows how the application can show the user how homogeneous are the pathways of different callers, and where the pathways separate and come together again.

The previous figures displayed data from a very simple hypothetical TLC system. The application, however, is meant to view more complicated graphs of actual systems. Actual TLC systems come in a variety of complexities. The spinal cord dysfunction project (TLC SCD) is an example of a fairly complicated TLC system.

TLC SCD is a TLC system for people with spinal cord dysfunction (either spinal cord injury or multiple sclerosis). The goal of TLC SCD is to reduce the incidence of pressure ulcers and depression, which are particularly prevalent in this population (8). To test the application with data from an actual TLC system, data from the TLC SCD project was used (Figure 10).

RESULTS OF EVALUATION

There were two main issues that were brought up during the evaluation. The first was that the graph of a complicated system appeared very busy and was difficult to

understand and follow an individual pathway. It was suggested to use some degree of abstraction to help users view the overall graph more easily.

This is an interesting suggestion. For many TLC systems, the call can be divided into modules that are relatively independent of each other. Showing a simplified abstraction of some modules, while showing the fully expanded graph of modules of interest would reduce the amount of information that needs to be displayed at once. This would result in some loss of ability to see how modules interact with each other, but might be worth it for the gain in clarity, so it would be nice to evaluate this in a future version of the application.

The second issue that was identified was that the application only showed the questions that patients were asked and their responses, not all interactions with the system. Sometimes the system will give the patient feedback about their previous answer before moving on to the next question – this feedback is not shown in the visualization application. Additionally, not all of the logic in the calls is based on the previous response. The rules which determine which question to ask next may depend on a previous answer in the current call, an answer from a previous call, or based on other information specific to the patient. This is not reflected in the visualization since it only shows how questions were answered in the given call – so callers with the same answer to a question might take different pathways.

It would be nice to have a way to identify the variables/conditions that directed which path was taken, especially so that users looking at the visualization could easier understand the calls as patients experience them. Including these in the visualization would provide a more complete representation of the calls, which would be especially useful for

users unfamiliar with the system. Unfortunately, this information is not part of the data that is captured during the TLC calls – this might, however, be worthwhile to collect in future systems to allow it to be included in the visualization.

CONCLUSION

This application provides a novel way of viewing TLC data, which will hopefully engage users – both TLC developers already familiar with TLC systems as well as other scientists – to help them to gain new insight into and understanding of the data and the way in which patients interact with the system. For TLC developers, this might allow them to formulate new hypotheses about the patients in their studies. Pathways can be compared of patients who showed greater benefit from the system to those who showed less benefit. Also, a fuller understanding of the patients' interactions with the system can help to guide the development of future systems to make them more effective and pleasing to the patients. For example, to reduce user burden, sparsely used sections of the graph can be explored to see if they should be pruned or deemphasized, or if they are important, revised to make them more frequently used.

For users unfamiliar with the TLC system, the visualization will help them to understand what was previously a black box. Additionally, this approach to medical interview data promises to generalize to other paradigms of the medical interview.

REFERENCES

1. Friedman RH, Stollerman JE, Mahoney DM, Rozenblyum L. The virtual visit: using telecommunications technology to take care of patients. *Journal of the American Medical Informatics Association*. 1997 Nov-Dec;4(6):413-25.
2. Rubin A, Migneault JP, Marks L, Goldstein E, Ludena K, Friedman RH. Automated telephone screening for problem drinking. *Journal on Studies of Alcohol*. 2006 May;67(3):454-7; *AMIA Annu Symp Proc*. 2003:1-5
3. Glanz K, Shigaki D, Farzanfar R, Pinto B, Kaplan B, Friedman RH. Participant reactions to a computerized telephone system for nutrition and exercise counseling. *Patient Education and Counseling*. 2003 Feb;49(2):157-63
4. Mooney KH, Beck SL, Friedman RH, Farzanfar R. Telephone-linked care for cancer symptom monitoring: a pilot study. *Cancer Practice*. 2002 May-Jun;10(3):147-54.
5. Delichatsios HK, Friedman RH, Glanz K, Tennstedt S, Smigelski C, Pinto BM, Kelley H, Gillman MW. Randomized trial of a "talking computer" to improve adults' eating habits. *American Journal of Health Promotion*. 2001 Mar-Apr;15(4):215-24.
6. Reas C, Fry B. Processing 1.0 <<http://processing.org/>>.
7. Kaufmann M, Wagner D, Eds. *Drawing Graphs: Methods and Models*. Berlin: Springer, 2001. 87-120.
8. Houlihan B, Paasche-Orlow M, Wierbicky J, Zazula J, Dicker K, Cuevas P, Williams S. Development of a novel automated telehealth system for people with spinal cord dysfunction. Poster presentation at Spinal Cord Injuries 2009 (March 19-21, 2009).

FIGURES

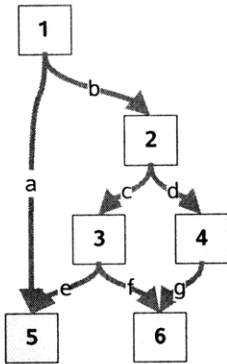


Figure 1. Directed graph drawn as Node-Link diagram. There are six nodes (1-6) and seven edges (a-g).

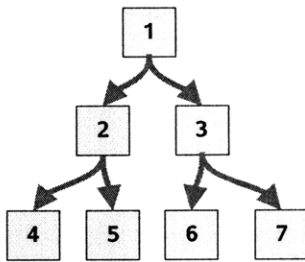
1. Do you have a cushion for the seat of your wheelchair? Please say yes or no.
If yes, GOTO 2
If no, GOTO 3

2. Does the cushion have gel inside, air inside or neither. Please say gel, air or neither.
If gel, GOTO 4
If air, GOTO 5

3. It's important to get a seat cushion for your wheelchair – would you like advice about how you can get one? Please say yes or no.
If yes, GOTO 6
If no, GOTO 7

...

a.



b.

Figure 2. (a) Sample TLC interview script excerpt (b) Script excerpt represented by a directed graph.

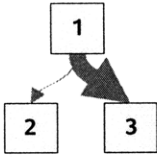


Figure 3. The thickness of the edge between nodes 1 and 3 (and relative thinness of the edge between nodes 1 and 2) indicates that patients are more frequently asked question 3 than question 2 after question 1.

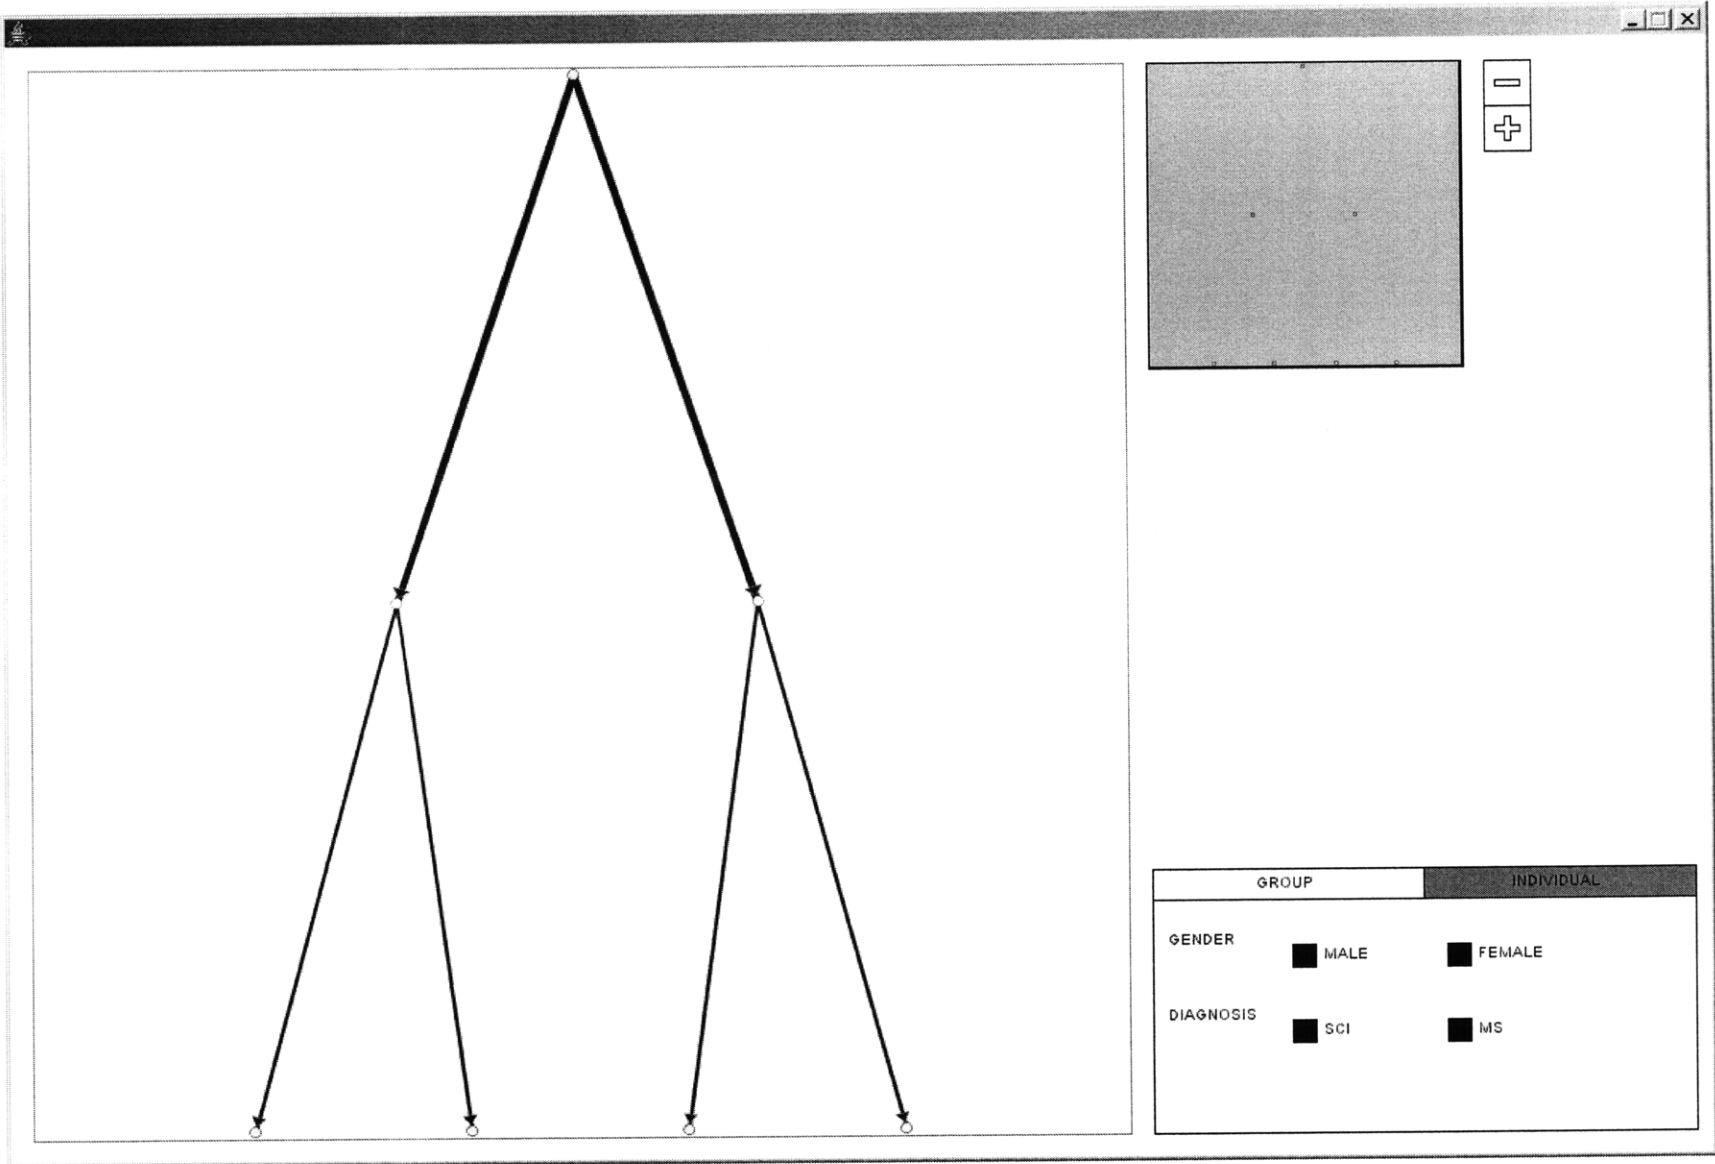
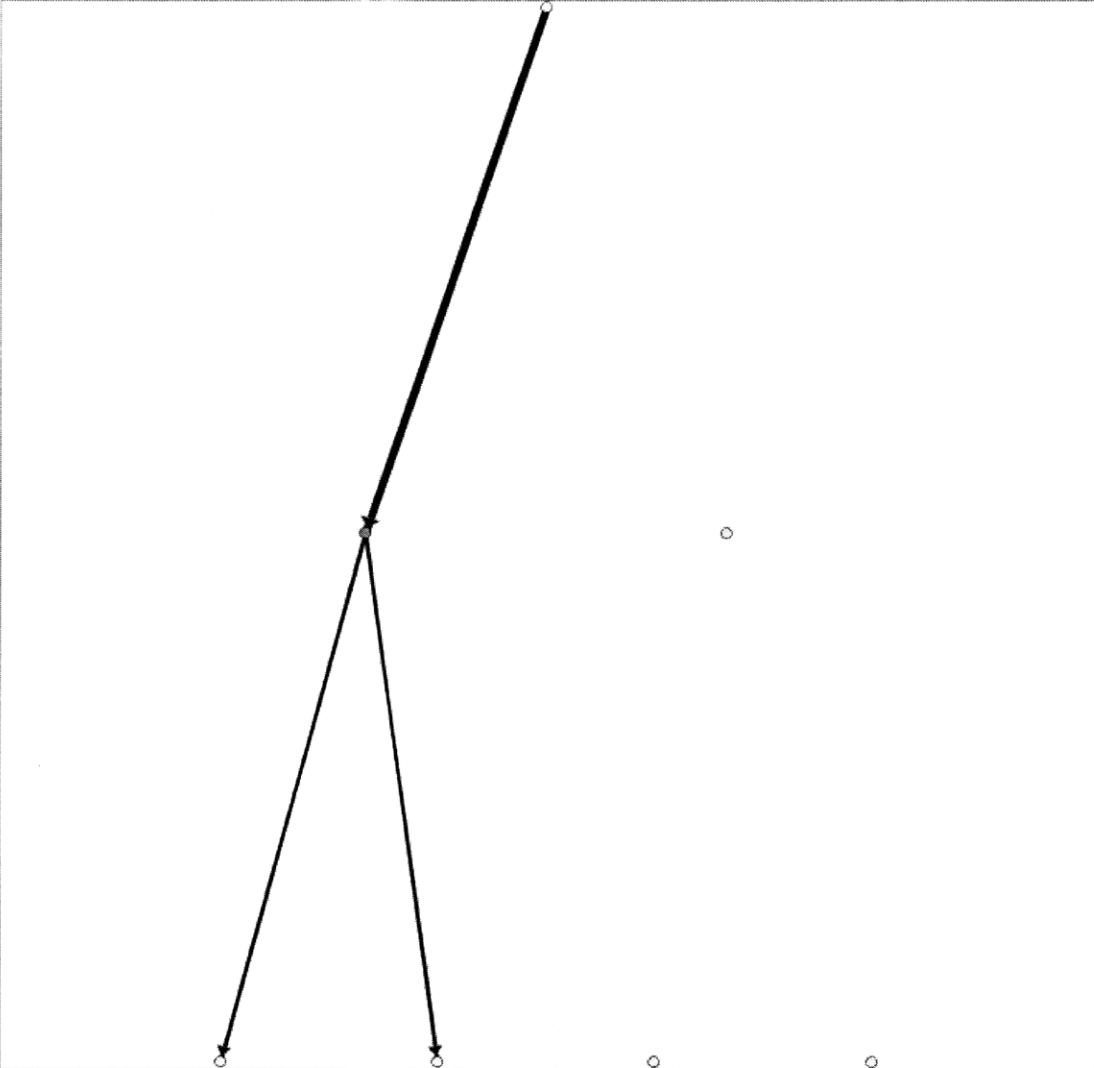


Figure 4. Application Window with data from sample TLC system described in Figure 2. The main window shows the seven question nodes (drawn as circles), and the edges connecting them. The small square on the right shows a smaller version of the question nodes, and the gray area shows which part of the graph is shown in the main window (in this case, the whole graph).



Does the cushion have gel inside air inside or neither? Please say gel air or neither.

GEL

AIR

GROUP		INDIVIDUAL	
GENDER	<input type="checkbox"/> MALE	<input type="checkbox"/> FEMALE	
DIAGNOSIS	<input type="checkbox"/> SCI	<input type="checkbox"/> MS	

Figure 5. Application Window with question node selected. When a question node is clicked on, it changes color and only the edges to and from that node are shown.

Figure 6. Application Window with one answer ("GEL") selected.

The screenshot shows a software application window with a dark title bar. The main area is divided into three sections:

- Decision Tree:** A tree diagram with a root node at the top. Two thick black lines connect the root to two intermediate nodes. From each intermediate node, two thin black lines connect to a total of four leaf nodes at the bottom.
- Image and Question:** A small grayscale image of a cushion is shown in the top right. Below it is the text: "Does the cushion have gel inside air inside or neither? Please say gel air or neither." To the right of the image are zoom-in (+) and zoom-out (-) buttons.
- Response Form:** A form with two rows: "GEL" and "AIR". Each row has a horizontal bar with a black fill, indicating that "GEL" is the selected answer. A small "1" is visible at the bottom right of the form.
- Attributes Panel:** A table at the bottom right with two columns: "GROUP" and "INDIVIDUAL".

	GROUP	INDIVIDUAL
GENDER	<input checked="" type="checkbox"/> MALE	<input checked="" type="checkbox"/> FEMALE
DIAGNOSIS	<input checked="" type="checkbox"/> SCI	<input checked="" type="checkbox"/> MS

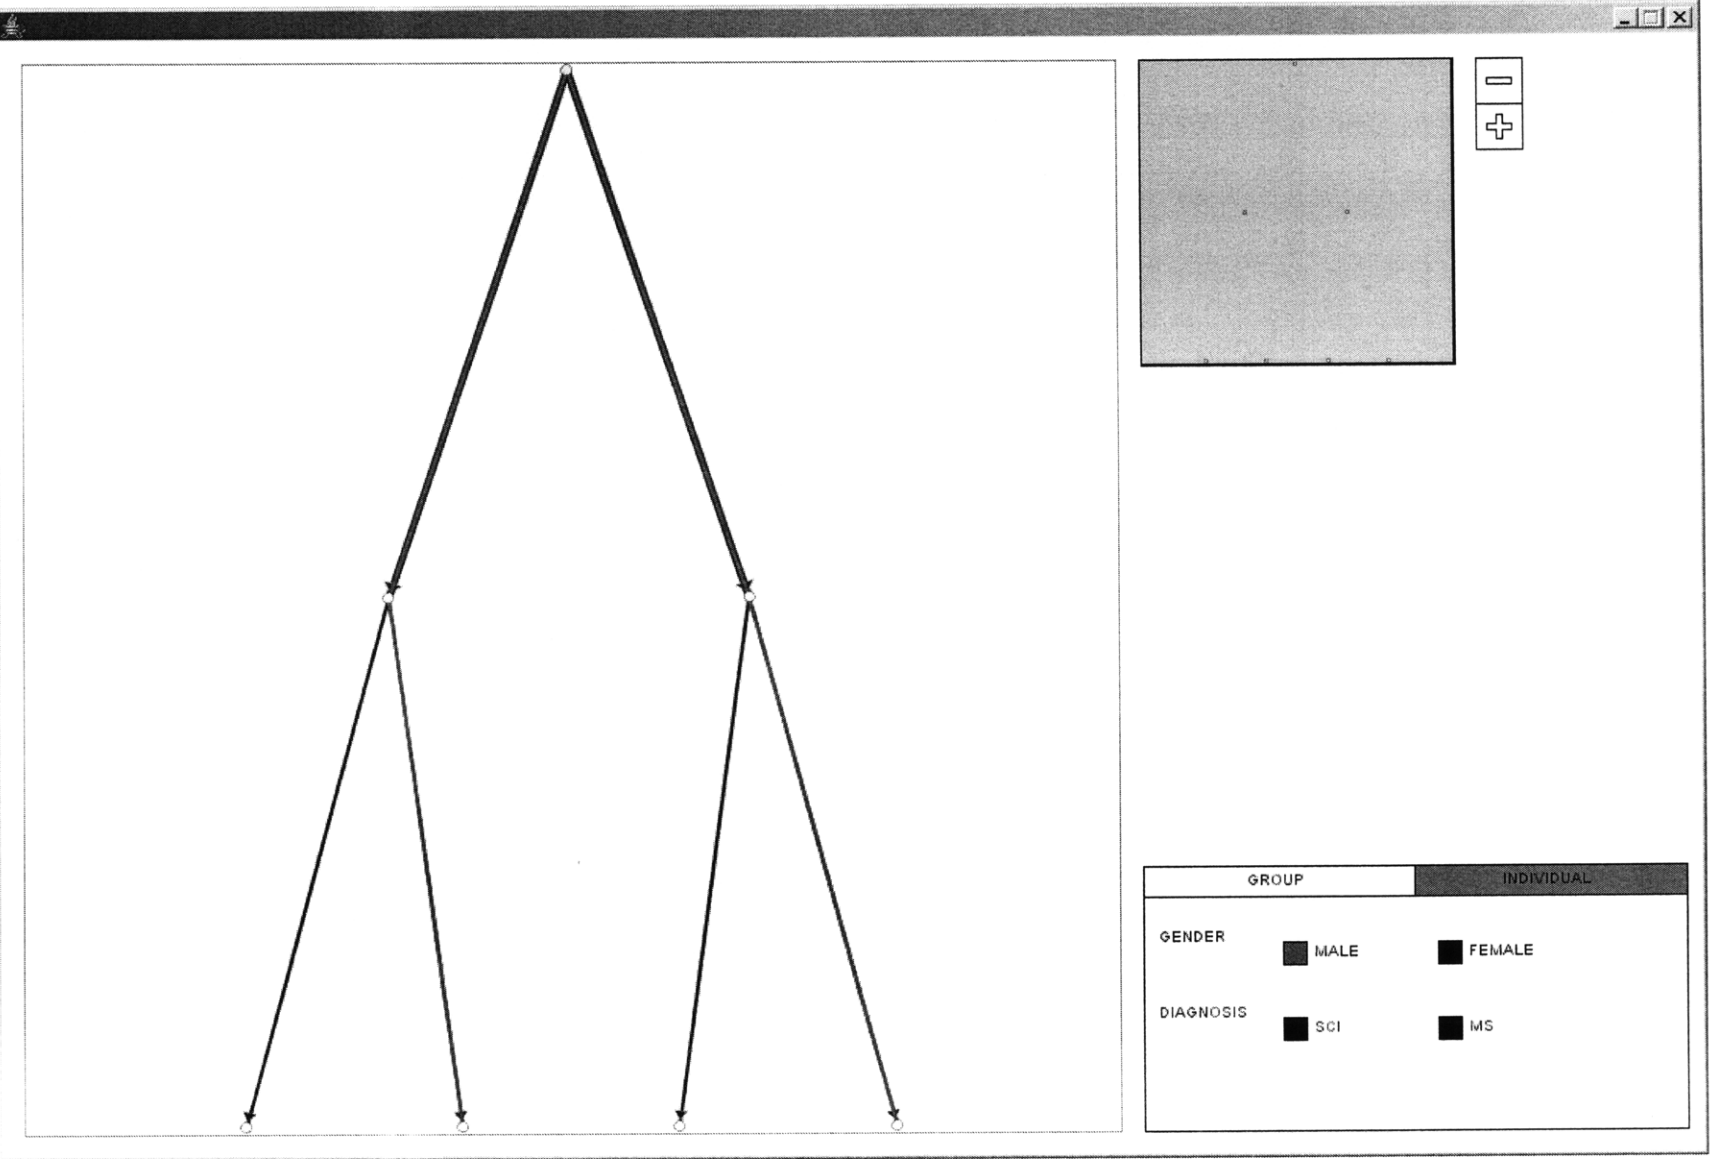


Figure 7. Application window with demographic type "MALE" selected.

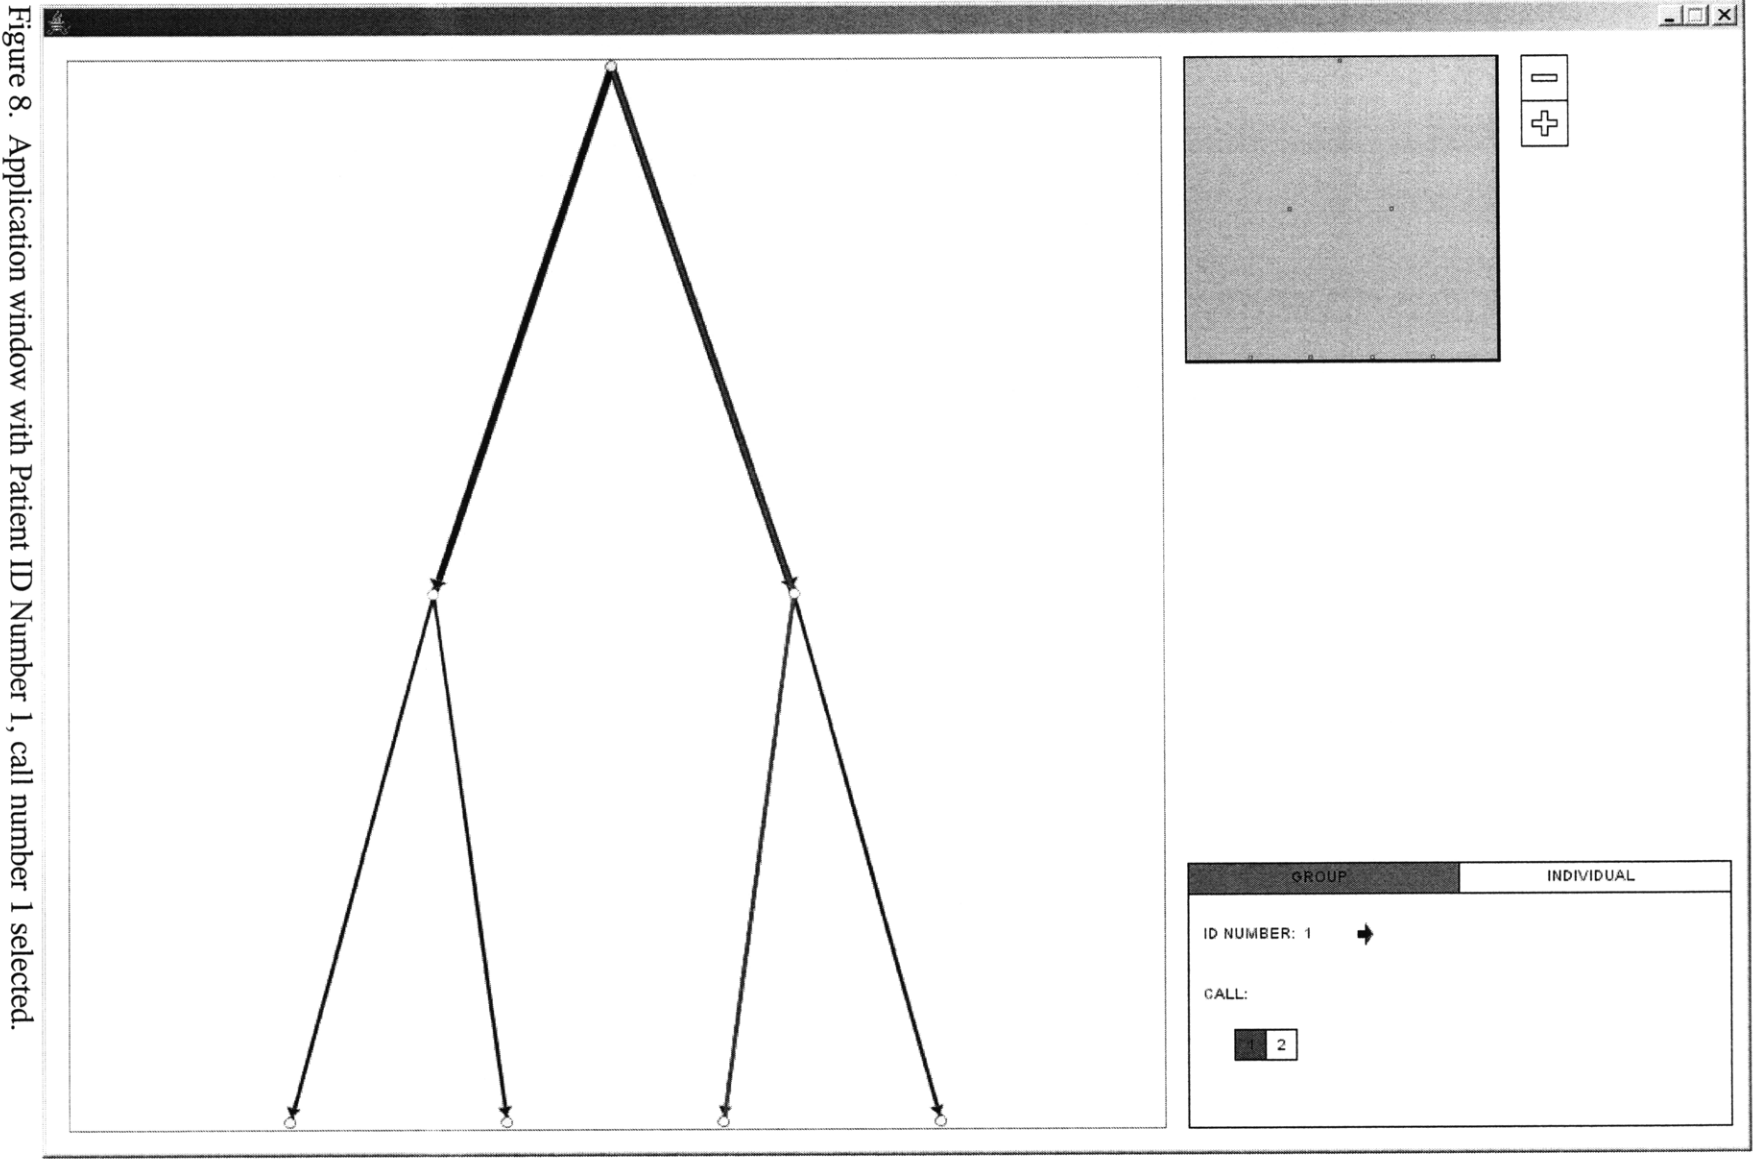


Figure 8. Application window with Patient ID Number 1, call number 1 selected.

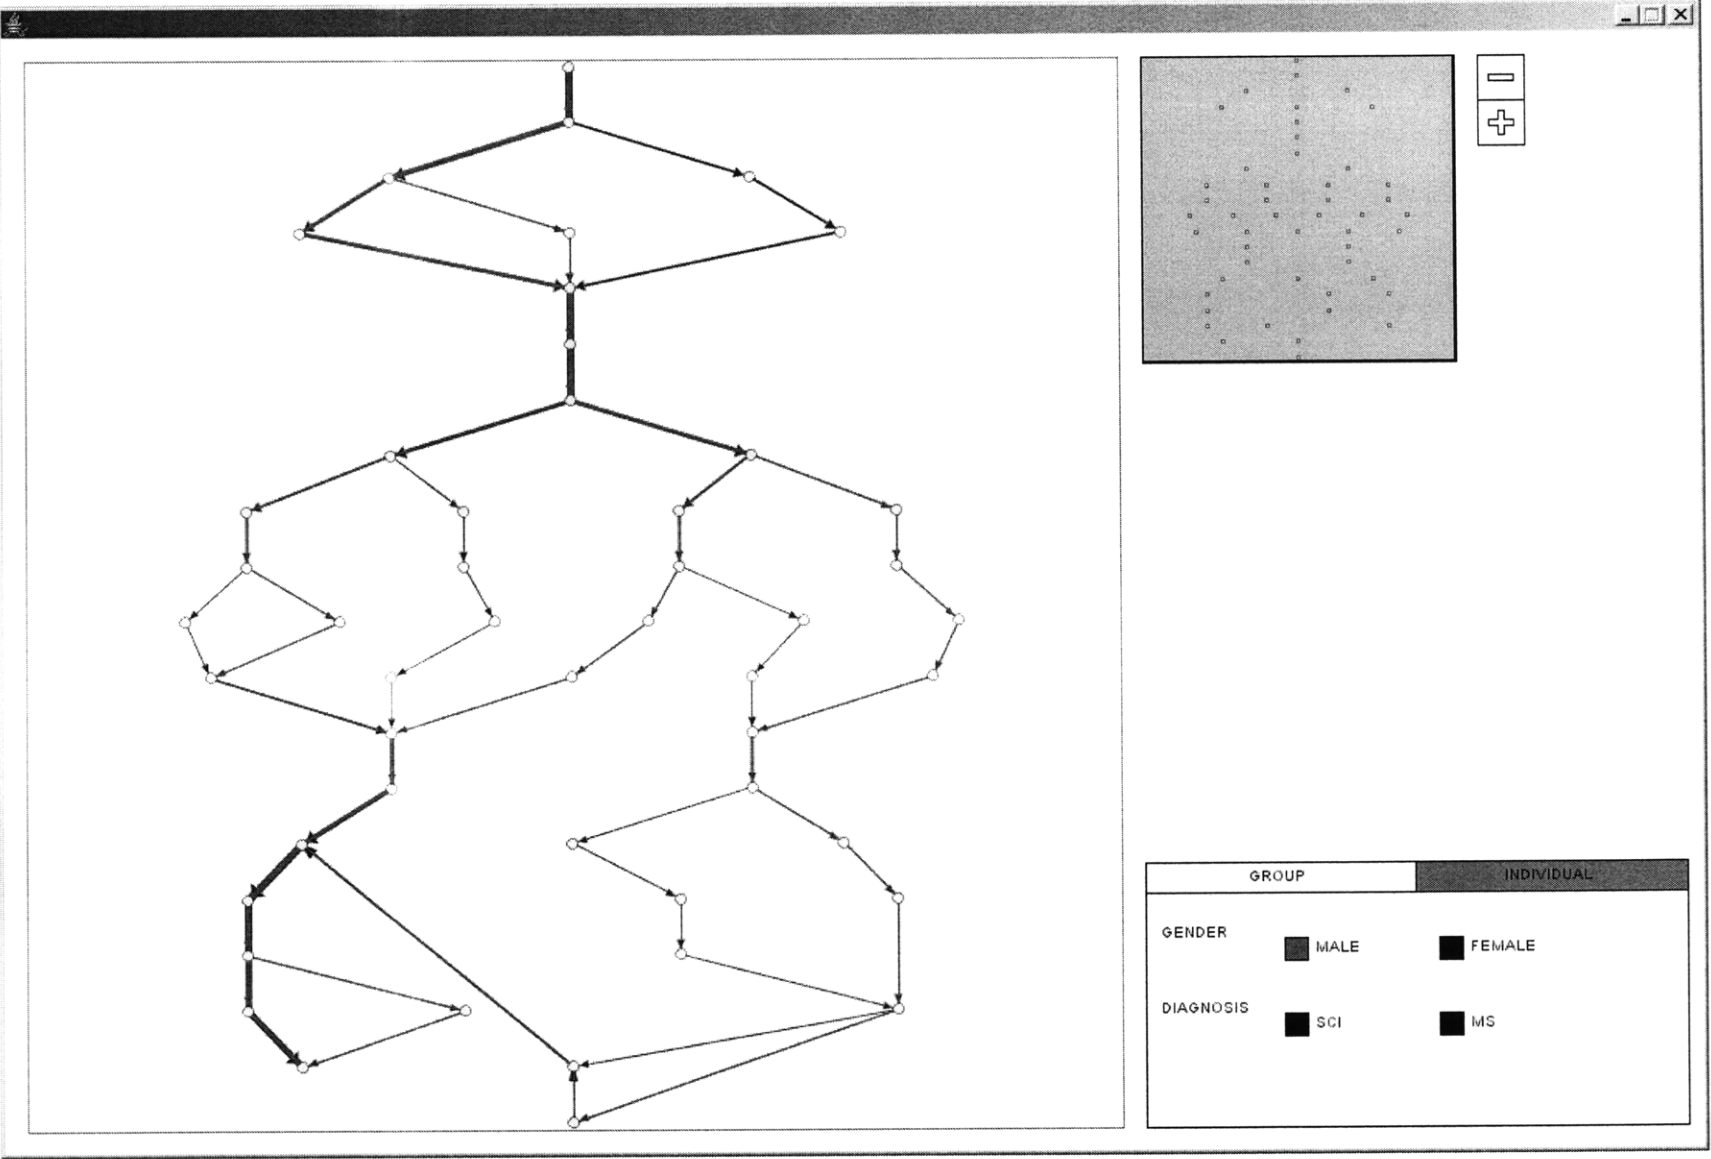


Figure 9. More complicated graph ("MALE" selected).

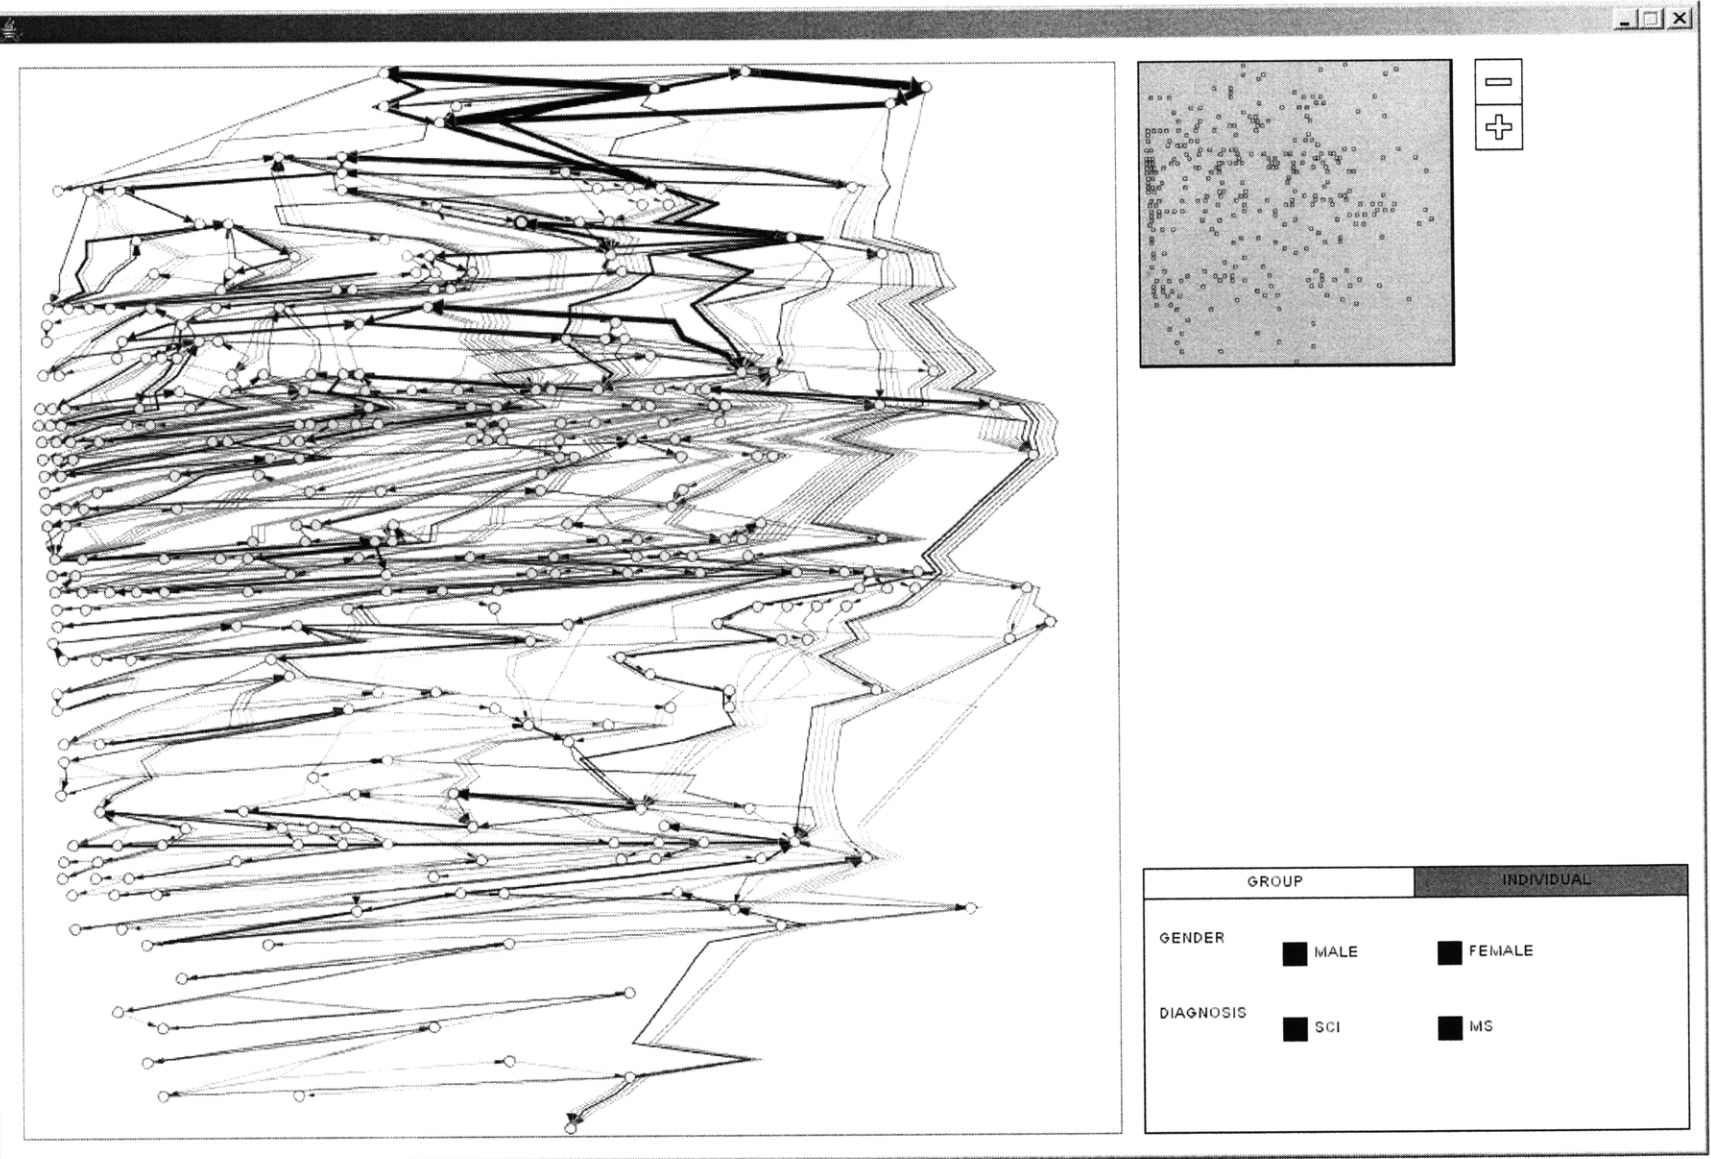


Figure 10. More complicated graph (data from TLC SCD project).