



American Society of  
Agricultural and Biological Engineers

[Go Back](#)

# Interactive Visualization of Multi-Dimensional Data in Dairy Production

D. Pietersma, J. P. H. Holthuis, R. Lacroix, K. M. Wade

---

Published in Applied Engineering in Agriculture Vol. 21(6): 1081-1088 (  
Copyright 2005 American Society of Agricultural Engineers ).

Article was submitted for review in August 2003; approved for publication by the Information & Electrical Technologies Division of ASAE in June 2005.

The authors are **Diederik Pietersma**, Postdoctoral Fellow, Dairy Information Systems Group, Department of Animal Science, McGill University, Macdonald Campus, Ste. Anne de Bellevue, Quebec, Canada; **Jurgen P. H. Holthuis**, Visiting Graduate Student, Department of Agricultural Engineering and Physics, Wageningen University, The Netherlands; **René Lacroix**, **ASABE Member**,

Assistant Professor, Dairy Information Systems Group, Department of Animal Science, McGill University, Macdonald Campus, Ste. Anne de Bellevue, Quebec, Canada; and **Kevin M. Wade**, **ASABE Member**,

Associate Professor, Dairy Information Systems Group, Department of Animal Science, McGill University, Macdonald Campus, Ste. Anne de Bellevue, Quebec, Canada. **Corresponding author:**

Kevin Wade, Dept. of Animal Science, McGill University, Macdonald Campus, 2111 Lakeshore Rd., Ste. Anne de Bellevue, QC, Canada H9X 3V9; phone: 514-398-7973; fax: 514-398-7964; e-mail: kevin.wade@mcgill.ca.

---

**Abstract.** *The objective of this study was to evaluate the potential usage, benefits, and limitations of using 3D graphics in an interactive exploration of multi-dimensional dairy data. A software prototype was, therefore, designed, consisting of two modules: a histogram to view joint frequency distributions for two dimensions of the data set at a time; and a scatter plot to explore clouds of data points plotted along three dimensions. A number of interaction techniques were implemented including chart zooming and translation, brushing of chart objects to display details, rapid attribute selection per chart coordinate, and dynamic data filtering. Chart rotation was implemented in order to analyze the 3D representations optimally. The feasibility of developing such prototypes was demonstrated in terms of data access, filtering, flexibility, and functionality under standard personal computer resources. However, performance tests with the 3D scatter plot revealed that in order for smooth interactions to occur, data needed to be limited to a few thousand records. The two resulting modules were seen as potentially useful in the support of various on-farm decision-making activities including the monitoring of cow*

*performance for the most recent milk test, and the benchmarking of herd performance, based on annual economics and milk recording data.*

**Keywords.** Interactive visualization, Data analysis, Software, Dairy farming.

In the dairy sector, large amounts of data are being collected and stored in databases. These data may yield valuable information to improve on-farm management. However, due to the large number of records and many attributes involved, discovering patterns and identifying abnormalities tend to be difficult. Dairy producers and their advisors may, therefore, benefit from using advanced software tools that help them explore the available data.

Visualization techniques have been developed in recent years to help users explore and explain large multi-dimensional data sets through interaction with visual representations (Card et al., 1999). Two different areas of visualization can be distinguished. Creating visual representations for data with an inherently 2D or 3D spatial structure is referred to as "scientific visualization," whereas designing appropriate visual representations for data that are not inherently spatial is called "information visualization" (Munzner, 2002). Data with a spatial structure are created with such applications as magnetic response imagery, computational fluid dynamics, and remote sensing (Grinstein and Ward, 2002). Inherently non-spatial or abstract data sets are often found in finance and other business applications (Wright, 1997). With information visualization, various spatial mappings are applied to abstract data, allowing for the exploitation of the human visual system to explore those data for decision-making purposes (Munzner, 2002). Data collected in the dairy sector tend to fall in the non-spatial category, involving measures of performance for individual dairy cows and aggregated performance indices for groups of cows and the entire herd. The application of information visualization techniques to dairy data sets might, therefore, support on-farm decision making.

Scatter plots in 2D are a very useful tool to detect patterns and outliers visually between two numerical attributes. To enable data exploration for more than two attributes simultaneously with such plots, the marker for each data point can be enhanced with identifiers (such as color, size, and shape) that vary with the value of the attribute they represent. However, attributes assigned to the X and the Y axes tend to dominate the user's perception of patterns compared to attributes assigned to aspects of the markers (Hoffman and Grinstein, 2002). In addition, when using color, extreme care should be taken to use a color map that accurately represents the data structure without introducing artifacts (Rogowitz and Treinish, 1996). Other techniques to enable visualization of multi-dimensional data sets with 2D graphs include the so-called parallel coordinates technique (Inselberg, 1985) and organized colored pixel displays (Keim and Kriegel, 1994).

Another approach to adding an additional attribute to the visualization is to make use of the 3D capabilities of modern graphics cards. Histograms representing joint-frequency distributions for two dimensions can be created with 3D graphics, while 3D scatter plots can give an overview of the shape of the data cloud created by three attributes and can help to find interactions among these attributes (Wilkinson, 1999; Ware, 2000). Techniques to provide a 3D effect on the 2D computer screen include perspective, occlusion (hiding of objects in the background by objects in the foreground), shading, motion, and stereoscopic viewing (Cleveland, 1993; Ware, 2000). However, 3D representations may create perceptual illusions and inhibit proper scale lookup to find data values, and occlusions may hide objects and patterns of interest (Shneiderman, 1996; Wilkinson, 1999). For these reasons, interaction methods like zoom and rotation may be required to enable effective visualization of 3D graphs.

Interaction techniques are not only useful for the support of proper depth perception with 3D graphs, they should also allow users to select and change attributes of interest and control the number of records shown, e.g. through data filtering techniques such as dynamic queries (Shneiderman, 1994). In addition to zoom and rotation, such techniques include brushing to show details of objects touched with the mouse pointer and distortion to show the area of interest with high detail while preserving an overview of the data (Keim, 2002). Interaction techniques

are thus an important part of the area of information visualization.

Commercial dairy herd management software increasingly includes graphs to support data analysis. However, these tools tend to include only standard 2D graphs and support only limited user interaction. Recent advances in computation power of processors and graphics cards, as well as the availability of third-party software components to support the creation of interactive 3D graphs, allow for the development of dairy information visualization applications. Extending dairy herd management software with visualization and interaction tools using 3D graphs might be helpful in the analysis of multi-dimensional dairy data. For example, such tools could be used by dairy producers and their advisors to find outliers among the cows in a herd (e.g., animals with a very low relative performance for several attributes of interest) and to compare a herd of interest with other herds for benchmarking purposes.

## Objectives

The overall objective of this research was to explore the use of interactive information visualization with 3D graphs to support dairy herd management. Specific objectives were to:

- design a 3D graph-based software prototype for interactive information visualization in dairy herd management,
- evaluate how the prototype could help with various decision-making tasks in dairy farming, and
- assess any limitations with 3D graph-based interactive information visualization for dairy herd management.

## Materials and Methods

### Interactive Visualization

Interactive visualization is a combination of visualization and direct manipulation techniques to improve the overall data exploration process. Interfaces that support direct manipulation (Shneiderman, 1983) have the following properties: 1) objects of interest are always visible; 2) users interact with the objects instead of using syntax; and 3) operations are incremental, rapid, and reversible. In addition, interactive visualization tools often make use of the so-called visual information-seeking mantra developed by Shneiderman (1996). This paradigm involves three steps: overview first, followed by zoom and filter, and finally, details-on-demand. The overview step shows users the entire data collection in one view. After this, users can zoom in on data items of interest and filter out uninteresting parts. With details-on-demand, users can select a data item or group of items to obtain detailed information.

### Data

Three data sets were used for software development, performance evaluations, and assessment of potential benefits of the application. The first data set involved test-day data for a single Quebec dairy herd covering seven milk test dates (84 to 94 lactating cows per test day) and was obtained from the Quebec dairy herd improvement agency (Programme d'analyse des troupeaux laitiers du Québec). On a test day – approximately once per month – the milk yield was measured for each lactating cow and milk samples were taken for subsequent lab analyses to determine milk constituents such as fat and protein percentages. Each record in this data set pertained to a specific test day and cow and included 15 attributes such as parity (i.e., number of calvings), body weight, days in milk after calving, milk, fat and protein yield, somatic cell count (SCC), and milk urea nitrogen (MUN). The two other data sets involved annual herd-average dairy

management data covering multiple years; the first consisted of mainly economic data for the years 1993 through 2000 for 821 to 1309 herds, and was obtained from an agricultural management and consultancy organization in Quebec (Fédération des groupes conseils agricoles du Québec). Each record was specific to a given year and herd. Of the 57 attributes in this data set, many dealt with variable costs per liter of milk produced (e.g., costs for purchased feeds, reproduction, and veterinary services), several cow performance indices (e.g., average milk yield, calving interval, and breedings per conception), and some attributes that described general herd characteristics (e.g., number of cows, type of breed, type of housing). The second of these data sets was obtained from the Quebec dairy herd improvement agency and comprised four years of annual herd average milk recording and nutrition data from 1999 through 2002 for 7011 to 6237 herds. This data set consisted of 43 attributes, such as herd average 305-day milk yield, fat and protein percentages, age and weight of the cows, amounts fed, and costs of various feeds.

Each data set was stored in an Access (Microsoft Corporation, Redmond, Wash.) database. Each database consisted of a single data table with two nominal fields to identify each subject (i.e., cow or herd) and the time related aspect of the measurements (i.e., date of test day or year), and numerical fields for the attributes of interest. Nominal attributes, such as type of breed, were numerated to enable plotting with the 3D graphs.

## Procedure

Software prototyping was used to develop a 3D graph-based interactive visualization application for the exploration of the three data sets. Two modules were developed: a histogram to view joint frequency distributions for two dimensions of the data set at a time; and a scatter plot to explore clouds of data points plotted along three dimensions. The visual information-seeking mantra (Shneiderman, 1996) was used as a guideline to ensure that a user first sees an overview of the data followed by the possibility of making adjustments to the data included, and viewing details of data items of interest. Interaction techniques that were implemented in the software included chart rotation and zooming, translation to move the chart view, brushing of chart objects to display details, rapid attribute selection per chart coordinate, and a dynamic query (Shneiderman, 1994) in the form of a data filtering slider.

The application was developed with Visual Basic version 6 (Microsoft Corporation, Redmond, Wash.). The interactive 3D graphs within the application were created with ComponentOne Chart version 7 (ComponentOne LCC, Pittsburgh, Pa).

## Evaluation

Potential usage of the two modules was explored in the context of two decision-making tasks in dairy farming. Firstly, the tools may help to monitor the performance of individual dairy cows based on the most recent test-day results. Secondly, they may aid in benchmarking the herd performance based on annual herd-average data.

Tests to evaluate the chart interaction performance were carried out for the scatter plot module, which was computationally more demanding than the histogram module. A large set of 3D data was created by randomly generating numbers for three attributes. The number of records used by the application was increased in a step-wise manner to determine the maximum number of records that could be used given unconstrained chart interaction (such as rotation) and severely delayed chart interaction. The tests were performed with a moderate and an advanced hardware system. The moderate system consisted of a 677-MHz Pentium 3 processor (Intel Corporation, Santa Clara, Calif.) with a 16-MB base-level graphics card (ATI Xpert 128 from ATI Technologies Inc., Markham, Ontario, Canada). The advanced system consisted of a 2.2-GHz Pentium 4 processor with a 64-MB professional graphics card (NVIDIA Quadro4 700XGL from NVIDIA Corp., Santa Clara, Calif.). Both systems used a monitor with a 1280- × 1024-pixel resolution.

# Results and Discussion

## Histogram

The histogram module uses a 3D bar chart to show the joint frequency distribution for two attributes. Figure 1 represents a screen capture of this module used with the test-day data set in which each record represents a cow in the herd. In this illustration, the attributes "Protein-to-fat ratio" and "Milk Urea Nitrogen" are categorized into bins and plotted on two axes that form a grid on the floor of the 3D chart. For each non-empty bin, a 3D bar is shown with its height representing the record count. In figure 1, most cows have a protein-to-fat ratio between 0.76 and 1.02 and a milk urea nitrogen value between 7.4 and 11.8 mg/dl. The bar chart represents data from a single test date that can be selected with the "Time" list box. Selecting an entry in this list box followed by pressing the arrow-up or arrow-down key on the keyboard allows users to scroll through the list and explore, in real time, how the frequency distributions vary with time. The attributes can also be changed with list boxes and the number of bins, used to determine the joint frequency distribution, can be adjusted. The results of changing a list box item are shown immediately in the graph, allowing users to explore a large number of attribute combinations quickly.

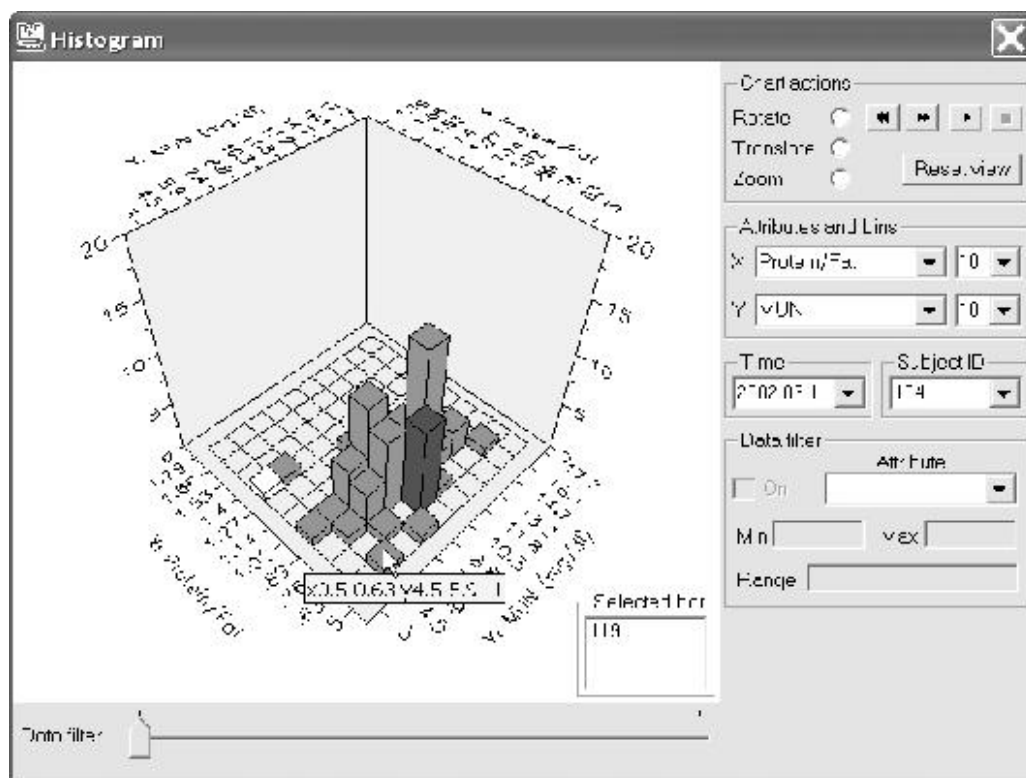


Figure 1. Screen capture of the histogram module showing the joint frequency distribution of protein-to-fat ratio and milk urea nitrogen for the cows in a herd on a particular test day.

In addition to selecting values for chart attributes and time, users have the ability to change the chart view by dragging the mouse pointer over the chart. Depending on the user-selected mode of action, this will result in a rotation, translation to move the area that is in focus, or a zoom. To support the use of chart rotation, buttons were added for constrained rotation around the Z axis. Clicking a "play" button provides a continuous animation, while holding down the fast forward or fast backward buttons allows users to rotate the chart to a desired view quickly. Moving the mouse pointer over a bar in the chart, i.e. brushing, displays a label with the bin ranges and frequency count of that bar (fig. 1). The position of a particular subject in the histogram can be found by selecting the identifier for that subject within a list box. For example, in figure 1, cow number 104 has been selected in the Subject ID list box thereby changing the color of the corresponding bar to which she belongs (dark gray bar representing a low protein-to-fat ratio

combined with medium value for milk urea nitrogen). Alternatively, the user can click on a bar to get a list of the subjects that belong to it. For example, in figure 1, the user has selected a bar (colored white) with a very high protein-to-fat ratio combined with a very low value for milk urea nitrogen. The "Selected bar" list box within the chart window reveals that this bar represents a single cow with ID number 119. A data filter was implemented to allow users to explore the effect of the level of a third attribute on the shape of the histogram. In figure 1, the data filter has not been activated; its functionality is explained in detail in the scatter plot section.

Users can easily perceive the intended depth of the 3D model rendered on the 2D screen, because the histogram bars are placed in perspective on the chart, with the objects in the foreground presented larger than the objects in the background. While the partial or full occlusion of objects in the background by objects in the foreground supports proper depth perception, tall bars in the foreground may hide interesting patterns in the background. Rotation is, therefore, an essential interaction technique to use with this module.

## Scatter Plot

The scatter plot module employs a 3D scatter chart and several interaction techniques. The user can select attributes of interest for each of the three axes that form the 3D space of the chart. Figure 2 represents a screen capture of this module used with the annual herd average economics data set, showing attributes "Milk revenue minus total variable costs," "Milk production per cow," and "Body weight per cow." Each data-set record is represented with a marker leading to a cloud of data points. In figure 2, body weight and milk per cow seem to be positively correlated, as expected, while the relation with milk revenue minus total variable costs is less clear. As with the histogram module, the chart represents data associated with a single time related item, i.e. "Year" for this data set. The user can scroll through the "Time" list box to evaluate changes in the shape of the visual representation and changes in the position of a selected subject within the data cloud.

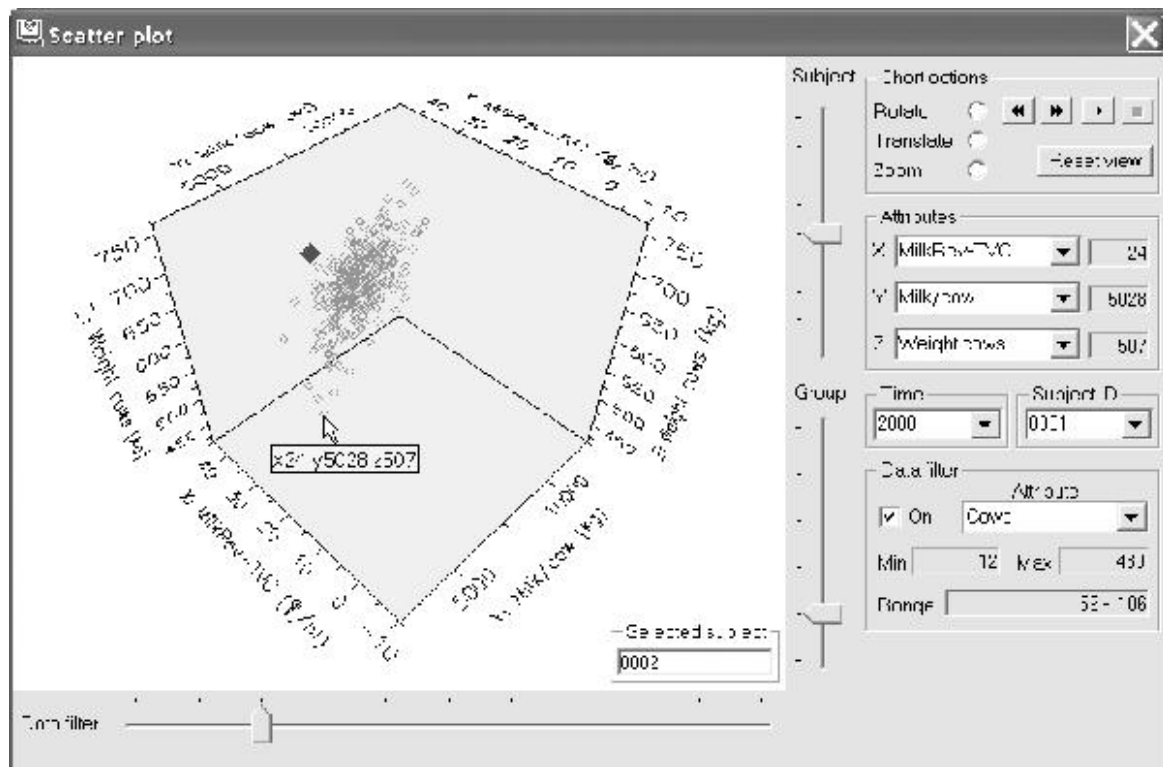


Figure 2. Screen capture of the scatter plot module showing the weight of cows plotted against milk production and milk revenue minus total variable costs, for the herds that are within the 59 to 106 cows per herd data filter range and for a particular year.

Users can interact with the visual representation in several ways: rotation, translation, and zoom

can be performed by dragging the mouse pointer over the chart, and buttons were added for constrained rotation around the Z axis. Brushing a data point in the chart (by moving the mouse pointer over it) reveals the attribute values (fig. 2). Clicking on a data point in the chart reveals the subject's identifier. Selecting a specific subject with a list box changes the marker representing that subject to a closed diamond of a different color and a larger size. To help the user find the selected subject in the data cloud, two vertical sliders (fig. 2) were added to adjust the marker size for either the selected subject or the non-selected data.

A horizontal slider control was added to allow users to view how the visual representation changes in 3D space in real time with increasing or decreasing levels of a filter attribute. Implementation of this control involved partitioning of the records in the data set into ten sub-sets or bins depending on the value of a filter attribute. For this purpose, the filter attribute range was subdivided into ten equal width intervals. For example, in figure 2, attribute "Number of cows in the herd" has been selected for data filtering and the slider has been moved to the right to the second level and bin, showing only the data of herds with 59 to 106 cows. Dragging the slider one notch further to the right would reveal the pattern representing herds in the next bin of the filter attribute. The number of slider settings and bins was automatically reduced when dealing with small data sets.

With the scatter plot module, the intended depth within the visual representation is much more difficult to perceive than with the histogram module. For static images, the scatter plot chart component provides only occlusion as a depth cue, while the histogram component included also perspective for each bar (fig. 1). However, with scatter plots, techniques like perspective and shading are expected to provide poor depth cues due to the small size of each marker (Ware, 2000). Stereoscopic viewing can help to perceive depth among data points that are close to each other, but requires specialized software and hardware. The so-called structure-from-motion cues created by chart rotation are especially helpful to perceive depth for data points that are far apart from each other (Ware, 2000). Thus, using chart rotation is very important with this module, not only to view data points hiding behind others, but also to help perceive the intended depth of the 3D data point cloud.

## Usage Scenarios

Potential usage of the two modules was explored for both monitoring of the performance of individual dairy cows using the most recent test-day results and benchmarking of the herd performance using annual herd average data.

### Monitoring

The monitoring of the performance of dairy cows following a milk test can be seen as part of operational-level or short-term decision making (Pietersma et al., 1998). For example, the most recent test results may reveal cows with poor performance or abnormal values for milk components. In addition, performance patterns may have changed compared to previous tests. Figure 3 shows three histograms for a series of milk tests. In figure 3a, representing the current test date, a particular cow has been selected with a very high value for attribute "Protein-to-fat ratio" and a very low value for attribute "MUN" (dark gray bar). Scrolling through the list box for time (fig. 1) reveals that the cow in question also had a low MUN value for the previous test (fig. 3b) and an extremely high protein-to-fat ratio for the test date prior to that (fig. 3c). The high protein-to-fat ratio and low values for MUN are indicative of nutritional problems for this particular cow.

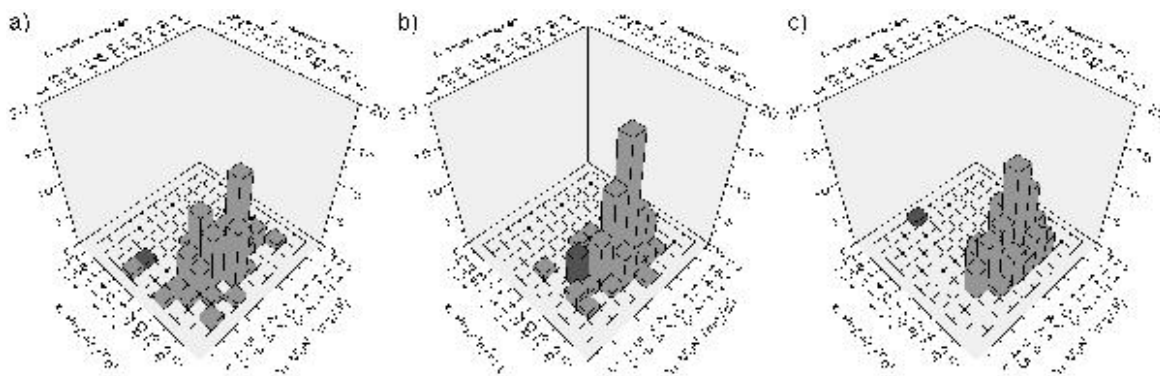


Figure 3. Histograms showing a selected cow within the joint frequency distributions of protein-to-fat ratio and milk urea nitrogen for the current and two previous milk tests.

The scatter plot module allows users to view the exact position of each cow within the herd instead of the discrete position in the histogram, and adds a third attribute to the visual representation. Figure 4a shows the scatter plot for the same data and selected cow as represented by the histogram of figure 3a, with the attribute "Protein percent" added along the Z axis. However, it is quite difficult to perceive the position of each point in the data cloud properly. Rotating the chart helps to perceive the 3D structure, as illustrated in the sequential plots of figure 4b and 4c. The scatter plot reveals in one view, supported by rotation, that the selected cow has a high protein-to-fat ratio combined with low MUN and medium protein. Thus, the scatter plot allows for the detection of outliers among three attributes at a time, while the histogram is limited to two.




Figure 4. Scatter plots showing a selected cow within views of protein plotted against protein-to-fat ratio and milk urea nitrogen (MUN) for three sequential rotation positions.

The slider control can be used to explore whether the pattern among protein-to-fat ratio, MUN, and protein – as shown in figure 4a – varies with increasing and decreasing values for a filter attribute such as "Days in milk." Figure 5 shows the result with each scatter plot representing a different position of the slider and associated range of filter attribute values. Extreme values for each of the three attributes seem to occur more often with the first slider setting, representing early lactation (4 to 101 days in milk), than with the next two slider settings representing mid- and late lactation.






Figure 5. Scatter plots showing views of protein plotted against protein-to-fat ratio and milk urea nitrogen (MUN) for three slider positions representing subsequent ranges for days in milk.

## Benchmarking

Benchmarking of a dairy herd involves comparison with reference herds to find management strengths and weaknesses. Benchmarking can be considered as part of tactical-level or medium-term decision making (Pietersma et al., 1998) and annual herd-average data are often used for this exercise. Figure 6 shows three charts of the histogram module for the most recent and two preceding years of dairy economics data in which a herd of interest has been selected. For the current year (fig. 6a), the selected herd had a fairly high value for the attribute "Milk revenue minus total variable costs" and a very low value for the attribute "Purchased feed costs." The histograms of the two preceding years (figs. 6b and 6c) reveal that the population of herds moved to higher values for milk revenue minus total variable costs and that the selected herd improved its relative position for both visualized attributes.

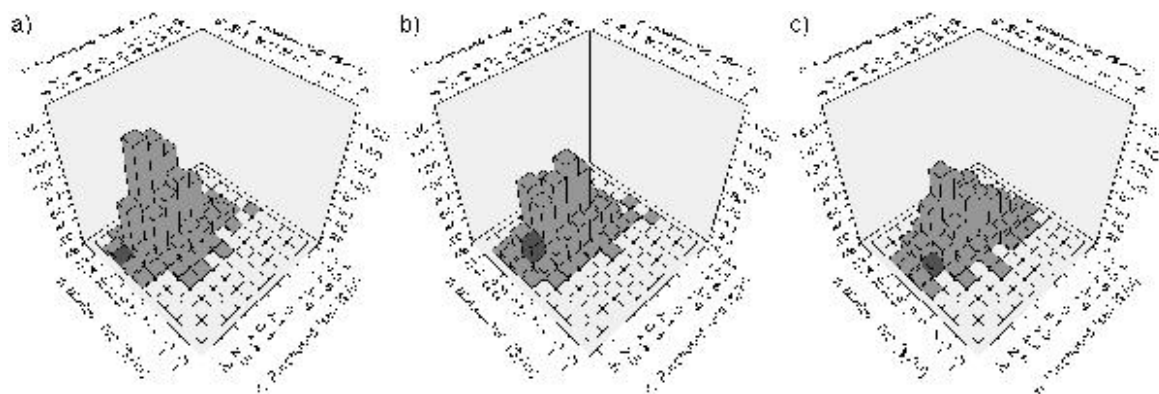


Figure 6. Histograms showing a selected herd within the joint frequency distributions of milk revenue minus total variable costs and purchased feed costs for most recent and two preceding years.

The annual herd average milk-recording and nutrition data set can also be used for benchmarking. While the large number of records (more than 6000 dairy herds) does not compromise the use of the histogram module, it does lead to some challenges when using the scatter plot tool. Figure 7a shows the histogram for the attributes "SCC" and "Calving interval," with the bar for the herd of interest highlighted. This herd has a SCC value similar to most other herds, while the calving interval is shorter than average. Plotting the same attributes together with the attribute "Breeding" in the scatter plot (fig. 7b) results in a data cloud with many points and a large amount of occlusions. In figure 7c the number of reference herds has been decreased with the data filter control to limit the view to herds with a similar average milk yield, making it easier to evaluate the position of the herd of interest within the 3D structure. Rotation is required to evaluate the value for the "Breeding" attribute, which, in this case, turns out to be average for the herd of interest.




Figure 7. Histogram showing a selected herd within the joint frequency distribution of somatic

cell count and calving interval and scatter plots showing the same selected herd within the data cloud for breedings plotted against somatic cell count and calving interval without and with data filtering for milk yield.

## Performance

Rotation requires a large amount of processing to calculate the position of each point on the 2D screen space and should occur in real time with smooth movements to allow users to perceive the 3D structure easily and without distractions. For the moderate hardware system, chart interaction without noticeable delays was limited to approximately 1500 data points. With the advanced system, uncompromised rotation was possible with up to approximately 5000 records. Given very long interaction delays, the moderate and advanced systems were able to handle up to 0.5 and 1 million data points, respectively. When substantial interaction delays occur, effective comparison of different chart views could be achieved by plotting a series of charts in parallel on the screen, but this would reduce the available surface area per chart. For the test day records ( $n = 94$ ) and the dairy economics data set ( $n = 1309$ ) used in this research, uncompromised chart interaction was possible on either hardware system. However, chart rotation with the 3D scatter plot module while exploring patterns among all 6237 herds in the milk-recording data set caused noticeable delays with the advanced system and severe delays of 1 to 2 s with the moderate hardware system. With more advanced hardware, uncompromised interaction with larger data sets would be possible. However, scatter plots of very large data sets tend to result in dense data clouds, thereby making it difficult to detect patterns of interest. In such situations, it may be more useful to use cubes that are positioned in a 3D grid and that vary in size to represent the joint frequency distribution of three attributes.

## Conclusion

An application for interactive visualization of dairy data was developed using a commercially available 3D chart component. The use of 3D graphs allowed for the investigation of complex patterns in the data, such as joint frequency distributions and interactions among three attributes at a time. Several user-interaction techniques were implemented, including rapid attribute selection, viewing of patterns in time, and filtering of the data depending on the level of an additional attribute. Due to the limitations of rendering on a 2D screen, chart interaction techniques, such as rotation, needed to be implemented to enable appropriate interpretation of the data. The 3D bar chart for joint frequency distributions of two attributes was considered easy to use. Proper perception of depth with the 3D scatter plot was considered difficult, requiring much use of the chart rotation functionality. Rotation of the 3D scatter plot without noticeable delays limited the size of the data set to a few thousand records. Interactive visualization with 3D graphs can be a useful addition to dairy herd management software to support the analysis of multi-dimensional data.

## Acknowledgements

The authors wish to acknowledge the financial support of the Fonds québécois de la recherche sur la nature et les technologies, Novalait inc., and the Ministère de l'Agriculture, des Pêcheries et de l'Alimentation du Québec. Data were provided by the Quebec federation of agricultural management and consultancy groups (Fédération des groupes conseils agricoles du Québec) and the Quebec Dairy Herd Analysis Service (Programme d'analyse des troupeaux laitiers du Québec).

## REFERENCES

Card, S., J. Mackinlay, and B. Shneiderman. 1999. *Readings in Information Visualization: Using*

*Vision to Think* . San Francisco, Calif.: Morgan Kaufmann.

Cleveland, W. S. 1993. *Visualizing Data* . Summit, N.J.: Hobart Press.

Grinstein, G. G., and M. O. Ward. 2002. Introduction to data visualization. In *Information Visualization in Data Mining and Knowledge Discovery* , eds. U. Fayyad, G. G. Grinstein, and A. Wierse, 21-45. San Francisco, Calif.: Morgan Kaufmann.

Hoffman, P. E., and G. G. Grinstein. 2002. A survey of visualizations for high-dimensional data mining. In *Information Visualization in Data Mining and Knowledge Discovery* , eds. U. Fayyad, G. G. Grinstein, and A. Wierse, 47-82. San Francisco, Calif.: Morgan Kaufmann.

Inselberg, A. 1985. The plane with parallel coordinates. *The Visual Computer* 1(2): 69-91.

Keim, D. A., and H.-P. Kriegel. 1994. VisDB: database exploration using multidimensional visualization . *IEEE Computer Graphics and Applications* 14(5): 40-49.

Keim, D. A. 2002. Information visualization and visual data mining . *IEEE Trans. on Visualizations and Computer Graphics* 8(1): 1-8.

Munzner, T. 2002. Information visualization. *IEEE Computer Graphics and Applications* 22(1): 20-21.

Pietersma, D., R. Lacroix, and K. M. Wade. 1998. A framework for the development of computerized management and control systems for use in dairy farming. *J. Dairy Sci.* 81(11): 2962-2972.

Rogowitz, B. E., and L. A. Treinish. 1996. Practical visualization: how not to lie with visualization. *Computers in Physics* 10(3): 268-273.

Shneiderman, B. 1983. Direct manipulation: a step beyond programming languages. *IEEE Computer* 16(8): 57-63.

Shneiderman, B. 1994. Dynamic queries for visual information seeking. *IEEE Software* 11(6): 70-77.

Shneiderman, B. 1996. The eyes have it: A task by data type taxonomy for information visualization. In *Proc. IEEE Symp. on Visual Languages* , 336-343. Piscataway, N.J.: IEEE.

Ware, C. 2000. *Information Visualization: Perception for Design* . San Diego, Calif.: Academic Press.

Wilkinson, L. 1999. *The Grammar of Graphics* . New York: Springer-Verlag.

Wright, W. 1997. Business visualization applications. *IEEE Computer Graphics and Applications* 17(4): 66-70.