

# An Evaluation of Interaction Techniques for the Exploration of 3D-Illustrations

Ian Pitt\*, Bernhard Preim\*\* & Stefan Schlechtweg\*\*

\*Department of Computer Science, University College Cork

\*\*Institut für Simulation und Graphik, Otto-von-Guericke-Universität

## Zusammenfassung

Dieser Beitrag beschreibt eine empirische Evaluierung von neuartigen Interaktionstechniken in einem hypermedialen Lehr- und Lernsystem, dem Zoom Illustrator. Die Integration von interaktiver 3d-Graphik mit Hypertextmöglichkeiten, die Nutzung von Fisheye-Techniken zur Navigation in textuellen Informationen und die Generierung von Bildunterschriften sind die Kernmerkmale des evaluierten Systems.

Der Zoom Illustrator ist besonders geeignet für die Anatomieausbildung, kann aber auch in anderen Bereichen zur Erklärung räumlicher Zusammenhänge genutzt werden. Die Interaktionstechniken des Zoom Illustrators wurden mit denen eines einfacheren Systems verglichen, das die selben Informationen benutzt und die gleiche Funktionalität bereitstellt. Das System wurde mit 8 Medizinstudenten bzw. jungen Ärzten getestet. Die Evaluierung zeigt die Nützlichkeit und die Akzeptanz der neuen Interaktionstechniken.

## Abstract

We present an empirical evaluation of innovative interaction techniques in an educational hypermedia system. The techniques investigated include the integration of 3d-interaction and hypertext-functionality, the use of fisheye techniques for the purpose of navigating in textual information and the generation of figure captions which automatically describe the images presented.

The system under evaluation is particularly useful for anatomy teaching, although it is not restricted to this domain. We compared a fully-fledged version of this system with a cut-down version which used only traditional interaction techniques which was otherwise identical. Both systems use exactly the same geometric models and textual explanations so that the comparison yields reliable results. Eight medical students and recently-qualified medical doctors took part in the evaluation. The results indicate the usefulness and the acceptance of the newly-developed interaction techniques.

## 1 Introduction

In many areas, learning involves the study of complex 3d-phenomena. In anatomy, for example, highly complex structures are studied under different aspects and viewing directions. To gain an insight into the spatial structure, and to be able to name parts of the objects, are important goals in the learning process. Similar problems arise in technical documentation where the construction of complex technical models must be communicated.

Traditional teaching materials do not support all aspects of this learning process sufficiently. Textbooks (see e.g. [1], [2]) contain valuable drawings but require a lot of flipping through the pages. Numerous illustrations and the text referring to them have to be mentally integrated. The Zoom Illustrator was developed in the light of these observations. It combines techniques from traditional media with the flexibility of a computer system. The system generates illustrations which are adapted to the parts currently being explained, thus allowing a detailed study of interesting parts while maintaining the context.

In this paper we describe a practical study in which we set out to evaluate the major features of the Zoom Illustrator and to establish the value of each in facilitating access to 3d-information.

## 2 The Zoom Illustrator: An Overview

The Zoom Illustrator and the techniques it embodies have been described in a number of publications, for example [3,4,5,6]. The following description is designed to give a brief overview, concentrating on those aspects of the design which are evaluated in this study.

The Zoom Illustrator is designed to permit the exploration of a 3d-model, displaying the image and its accompanying text within a single window. The user can then scale and rotate the model freely, whereupon the position and size of the text will be adjusted to preserve the context. A typical Zoom Illustrator display is shown in Figure 1 (next page).

Images and texts are located in different areas and connected to each other via reference lines. When the 3d-model is rotated and scaled, the reference lines are updated. To ensure that the image is not occluded by textual information and vice versa, the window is subdivided into a central part for the image (which occupies 50% of the space), as well as a left and a right part for textual descriptions (which occupy 25% each of the screen space available, see Figure 1).

The system is based on polygonal surface models of moderate complexity (~10,000 polygons) which are commercially available. These models are large enough to be interesting but small enough to be handled at interactive rates.

### 2.1 Use of Fisheye Techniques

Textual information includes labels and more or less detailed explanations which are connected to each other via hyper-links. The explanations were taken from well-established textbooks. Textual information is embedded in rectangular areas. The placement and size of these rectangles is based on fisheye techniques (as introduced in [7]). Other systems, by contrast, open a separate window to display an explanation, making it difficult for the user to create a mental link between the explanation and the object to which it refers. Fisheye techniques provide a useful way to zoom in and out interactively, presenting several levels of detail simultaneously. This makes it possible to explain objects in which the user is interested, while automatically scaling down and repositioning others to provide the necessary space without overlapping information.

The placement and the size of information is based on the degree of interest (DOI), an application-specific value influenced by user interactions. Continuously enlarging entities at the expense of others, the size of which are reduced accordingly, is accomplished by a variant of the continuous-zoom algorithm as presented in [8]. Whenever a rectangle becomes large enough to accommodate more text, an additional explanation is displayed which is subsequently hidden automatically if other rectangles are enlarged.

Objects which are textually explained are emphasized within the 3d-model. This strategy – although useful in many cases – is not adequate to illustrate small details. The illustration of small details is accomplished with the 3d-fisheye zoom [3,9]. It makes it possible to scale up small details while automatically repositioning other objects — the same strategy which has been successfully applied to text navigation. The application of fisheye techniques for both symbolic and graphical exploration provides a coherent interface.

## 2.2 Enhancing Navigation in Textual Information

One important question when applying zoom techniques concerns the automatic hiding of nodes. While the zoom algorithm generally works in a comprehensible manner, it could be irritating if one node disappeared because another node had grown in size. Even if users understand what has happened they might not know what they are supposed to do to get a node back. This raises two questions:

- How to prevent nodes from being closed?
- How to get nodes back which have been closed?

We developed one solution for each question.

To prevent nodes from being closed, we introduce an additional area, a *pinboard*, as a container for some privileged nodes (up to 4) which are not exposed to the zoom algorithm and therefore cannot be closed.

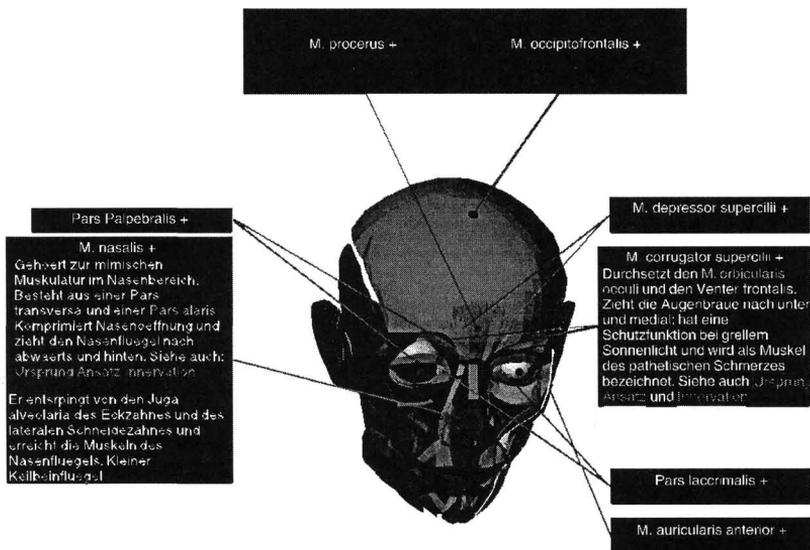


Figure 1: A typical Zoom Illustrator screen display. Text and images are displayed in a single window, with the image in the centre and the text in two side-panels, connected to points on the image by reference lines. More information about an object can be obtained by zooming its label. Surrounding labels will shrink to make space but remain visible. In the upper part is the pinboard, an area controlled by the user. In this example the user has attached two nodes to the pinboard.

The user can initiate an animated movement of a node to the pinboard (above the image) where it remains at a fixed position (see Figure 1). Nodes residing on the pinboard are still connected to the image via reference lines, which are updated if the 3D-model is transformed.

To get a node back once it has disappeared, it should be possible to select it via its label. For this purpose we designed a 3d-widget we call a *roundel* which contains all labels grouped

according to their category (see Figure 2). A 3d-widget fits naturally into our 3d-illustration system. The design of this widget is inspired by the work on 3d-interfaces carried out at XEROX Parc (see [10]). The roundel can be rotated by clicking on the disks at the lower and the upper part, one for a rotation to the left and one for the rotation to the right. The colour of the nodes which are closed is a saturated blue (instead of a weak grey for the nodes already presented), to encourage the user to invoke the node. The design and use of the pinboard and the roundel are discussed in [3].

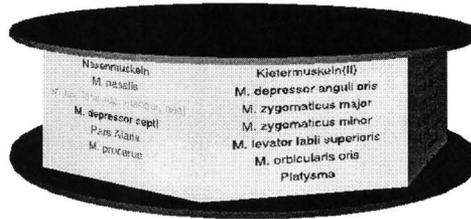


Figure 2: A roundel is used as an overview about all labels available. When a label is selected which does not belong to the illustration it is integrated in the illustration at an appropriate position.

### 2.3 Incorporation of Figure Captions

The Zoom Illustrator presents models which may be unfamiliar to the user, e.g. anatomic models of inner organs. The presentation of these models is not only a straight-forward rendering but includes an adaptation to the context. Thus objects which are textually explained are emphasized (coloured or enlarged if they are too small). Objects which are not important in a certain dialogue context are deaccentuated by being rendered semi-transparent so that other objects become visible.

All these modifications may remain unnoticed. Therefore, we generate figure captions which automatically describe an image. These captions are kept consistent with the image to which they refer (see Figure 3 and [6,12] for a description of the text-generation methods used). The structure, content and linguistic realisation of these captions are, again, strongly influenced by those in text-books. Figure captions can be adapted using a configuration dialogue to tailor them to the individual user, for example, the user can indicate the objects in which he or she is particularly interested, causing these objects to be monitored by the captions. The incorporation of figure captions is a distinctive concept of the Zoom Illustrator.

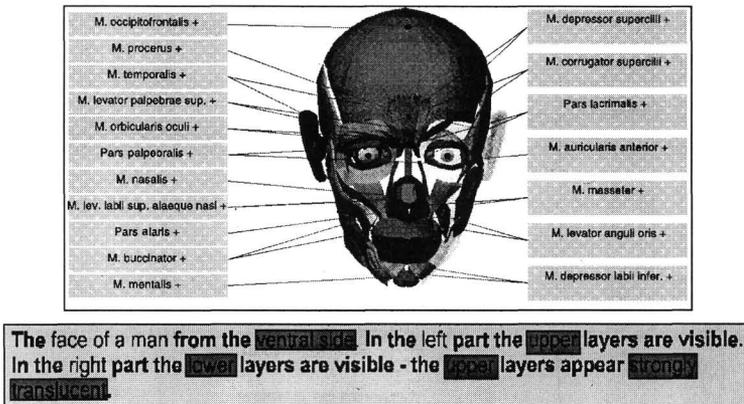


Figure 3: A figure caption describes the model depicted and important aspects of the image generation. In this caption, the difference between the two halves of a symmetric model is described. The dark rectangles correspond to sensitive parts of the caption. The selection of these parts leads to the presentation of pop-up menus with alternatives. Thus, the image generation can be influenced via the caption.

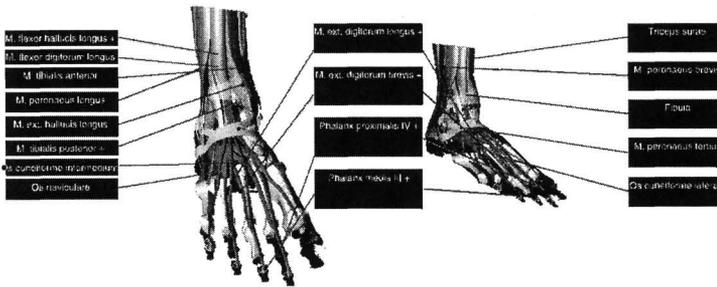


Figure 4: Two instances of a geometric model can be handled independently. Labels which refer to objects visible in both instances are placed between them to support the mental integration of both views.

### 2.4 Independent Handling of Two Models

A special feature of the Zoom Illustrator is its ability to display two instances of a 3d-model simultaneously. This results in an illustration with two models which can be manipulated independently from each other (see Figure 4). Labels which refer to objects which are visible in both instances are placed between them and connected to both 3d-models.

## 3 Evaluation

In this evaluation we set out to study the effectiveness of the distinguishing features of the Zoom Illustrator. In order to do this we decided to compare the Zoom Illustrator with another system which is similar in all important respects but uses only conventional interaction techniques. Since we wished to compare a number of features in single study we decided to use a questionnaire. A number of subjects were asked to perform a task using both the Zoom

Illustrator and the comparison system and then fill out a questionnaire in which they should rate the performance of the two systems on a number of scales. A within-subjects design was used because it was felt this would be sensitive enough to reveal differences between the two systems even with a relatively small sample size. In addition, we video-taped a few subjects to get a better impression about the interaction styles they prefer and the frequency of the use of certain features.

### 3.1 Choice of Comparison System

The choice of comparison system posed a number of problems. In order to obtain useful results it was important to compare the Zoom Illustrator with a system which is similar in all major respects save for those under test. In particular, the comparison system should provide the same level of detail in its models and graphical presentation and have the same basic interactive features. If the comparison system differed significantly in any of these respects it would be impossible to determine reliably the cause of any differences observed in the study.

A number of possible comparison systems were considered and rejected. One obvious approach was to compare the Zoom Illustrator with the VoxelMan (see [12]), a leading, commercially-available interactive anatomy trainer developed at the University of Hamburg. This would have offered a number of advantages, not least that it would have allowed us to compare the Zoom Illustrator with a well-known and widely-used system. However, the VoxelMan offers highly-detailed voxel models which differ significantly from the models used with the Zoom Illustrator. If VoxelMan were used as a comparison system, it would be very difficult to be certain that any differences observed did not stem from the disparity in level of detail available rather than from the different interactive facilities provided. Moreover, the types of model used in the two systems are not compatible, and a considerable amount of work would be necessary to convert between the volume models used by the VoxelMan and the surface models used by the Zoom Illustrator. Thus it is not feasible to compare the systems using a common model.

Another possibility was to compare the Zoom Illustrator with a CD-ROM version of an anatomy textbook (several of which are available). This would allow us to compare the Zoom Illustrator against a known system which has a fairly standard set of interactive features. However, CD-ROMs typically present highly-detailed (static) illustrations which contain far more detail than is available from the rendered images of the Zoom Illustrator. Thus, as with the VoxelMan, it would be very difficult to be certain that any differences observed were not caused by disparities in the level of detail available rather than from the different interactive facilities provided.

Having considered these and other options, it was decided that the Zoom Illustrator should be compared with a program written specifically for the purpose of the evaluation. This avoids the problem of differences in the level of detail available in the illustrations since the comparison system uses the same models as the Zoom Illustrator. The program mimics the operation and look of the Zoom Illustrator in all respects, except that the special interactive features under examination here are excluded. It is implemented with the same toolkit, namely Open Inventor, and the models are presented within the same 3d viewer (see Figure 5). The only differences are that the comparison system allows only one illustration to be viewed at a time, it does not generate figure captions, the explanations are presented in a separate window because no fisheye views are employed, and the user has to place labels manually.

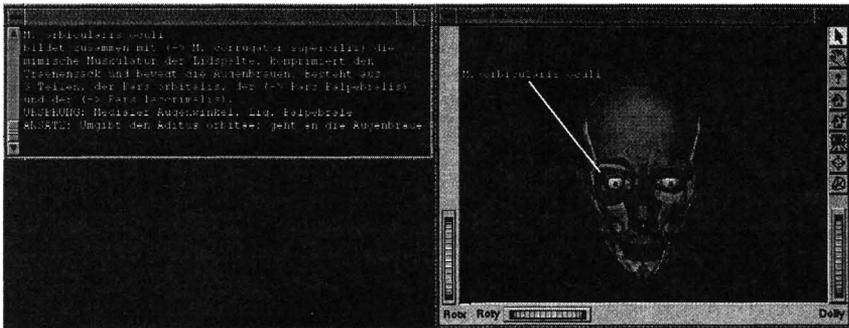


Figure 5: A typical situation working with the comparison system: The user has manually placed a label and asks for an explanation which is subsequently placed in a separate window.

### 3.2 Design of the Questionnaire

Having chosen a comparison system, we also had to design the questionnaire and decide which criteria to use when comparing the two systems.

Ideally we would have liked to have asked subjects to complete two questionnaires, one for the comparison system and one for the Zoom Illustrator. Subjects would have filled in the relevant questionnaire immediately after working with each system, and we would have obtained two sets of scores which could then be compared statistically.

However, this approach presented a major problem: it was very difficult to devise questions concerning the usefulness of features in the Zoom Illustrator which could be answered sensibly by a subject who was currently using the comparison system but who had not yet used the Zoom Illustrator.

For example, it would make little sense to ask a subject to rate the usefulness of, say, the pinboard, when that subject had not yet seen and used the pinboard. One possible solution would have been to ask more general questions concerning usability, maintenance of context, etc., but this does not fit well with the aim of this study which was to test the specific features of the Zoom Illustrator. Another possibility would have been to test all subjects on the Zoom Illustrator first and then on the comparison system, thus ensuring that when they came to fill in the questionnaires they had already used all the features mentioned. However, this might have biased the results since subjects often show a preference for the first system with which they are presented, and in consequence it is considered good practice to vary the order of presentation.

For these reasons, it was decided that subjects would be asked to work with both systems (half using the Zoom Illustrator first and half using the comparison system first) and then fill in a single questionnaire. For each feature, they would be asked to enter a value from 0 to 10, where 10 indicated that the feature was regarded as essential and should definitely be included (as in the Zoom Illustrator), and 0 indicated that the feature was of no value and should be excluded (as in the comparison system). A score of 5 indicated no preference in either direction. Using a single questionnaire in this way makes it more difficult to draw watertight conclusions, but avoids many of the problems inherent in testing features of a user interface.

The scales used in the questionnaire were chosen to reflect what we see to be the important features of the Zoom Illustrator compared with existing systems for viewing 3d-models, namely:

- The ability to view two instances of a model simultaneously
- The use of figure captions
- The use of fisheye techniques
- Automatic placement of labels

The first of these features, the ability to view two instances of a model simultaneously, can be tested on a single scale, as can the use of figure captions. Evaluating the use of fisheye techniques and the placement of labels is more complex, because a number of separate features have been incorporated into the Zoom Illustrator in order to implement fisheye techniques as effectively as possible. For example, the ability to shrink labels is part of the fisheye mechanism, but a further refinement in the Zoom Illustrator is that labels can be shrunk beyond the point of legibility rather than simply vanishing, thus providing a visual reminder of their position and content. Since these are separate features it was felt that they should be investigated separately in the evaluation.

The desire to implement fisheye techniques effectively is also reflected in a number of features which address some of the known problems associated with fisheye techniques. These include the problem of disorientation, which is addressed in the Zoom Illustrator through the use of fixed reference points and user-controlled information spaces (the pinboard and the roundel), and the problem of maintaining a clear link between images and text, which is addressed by ensuring that images and related text always change in tandem and by allowing users to position text within images. It was felt that the usefulness of each of these features should be examined separately.

We were also keen to discover the subjects' views on the use of hypertext links to move between views. These provide an obvious mechanism for linking visible objects to related objects in other parts of the model, but they may also disorientate the user.

The list of features included in the questionnaire was as shown below.

- Ability to view two illustrations at once
- Ability to change the size of the labels
- Ability to shrink labels without hiding/removing them
- The pinboard
- The overview offered by the roundel
- The integration of labels in helping users to understand the illustrations
- The ability to place labels inside an illustration
- Figure captions
- Hypertext functionality

In addition to the nine scales listed above, we introduced two further scales. One was designed to find out whether subjects felt the use of 3d-models was justified, or whether they felt that the spatial relationships could have been conveyed adequately using static images. The other scale concerned the quality of the models. The use of surface models of moderate

complexity allows the Zoom Illustrator to produce images in real time, but some might argue that this advantage is outweighed by the resultant loss of image quality. We wanted to find out whether the subjects felt the models had been of adequate quality for the task.

These last two scales differed from the others in that the subjects were asked to indicate their level of agreement with a statement. The possible responses ranged from 0 ('disagree') to 10 ('agree') with a response of 5 indicating no strong feelings in either direction. The two statements were:

- The 3d-presentation aids understanding more than would a 2d presentation.
- The quality and faithfulness of the models is appropriate for teaching topological relations.

### 3.3 Procedure

Eight subjects took part in the study. All had a basic knowledge of anatomy, being either medical students or probationary (i.e. recently-qualified) doctors. The subjects ranged in age from 22 to 31 and comprised 4 women and 4 men.

Each subject was first given a demonstration of the two systems. The demonstrations lasted for around 5-10 minutes (depending upon the number of questions asked by the subject, etc.) and the order of the demonstrations was varied so that half the subjects were shown the Zoom Illustrator first and half were shown the comparison system first. When the demonstrations were complete and any resulting questions had been answered, the subjects were asked to begin the main part of the experiment. They were asked to perform a simple task using first one system and then the other. The task involved exploring the *M. extensor digitorum longus*, a long and branching muscle, and correctly locating and tracing all its branches. The time required for completion of the task was typically around fifteen minutes and the order was varied so that half the students performed the task using the Zoom Illustrator first and half performed the task using the comparison system first. A few subjects agreed to be video-taped whilst carrying out the tasks. When they had finished, each subject was asked to complete the questionnaire.

### 3.4 Results

The results are shown in Table 1. It can be seen that in almost every case the average score is between 5 and 10, indicating that all the Zoom Illustrator's features were regarded as useful. In some cases (e.g., the ability to view two illustrations at once) the score is only slightly greater than 5, but in most cases it is considerably greater than 5.

The only exception to this pattern is the final scale, representing the subjects views on the quality of the models. The average score on this scale was below 5, reflecting a dissatisfaction with the models which was also expressed in informal comments by the subjects.

	S1	S2	S3	S4	S5	S6	S7	S8	Averages	Standard deviation
<b>Ability to view two illustrations at once</b>	10	8	1	3	5	10	4	8	6.13	3.36
<b>Ability to change the size of the labels</b>	7	10	8	10	7	10	10	7	8.63	1.51
<b>Labels shrink without vanishing</b>	7	10	8	8	10	5	10	9	8.38	1.77
<b>The Pin-Wall</b>	8	7	8	5	10	3	10	8	7.38	2.39
<b>The overview offered by the roundel</b>	8	10	8	8	8	5	10	9	8.25	1.58
<b>Integration of labels</b>	10	10	10	10	5	10	2	5	7.75	3.24
<b>Ability to place labels inside illustrations</b>	6	10	8	8	7	9	9	8	8.13	1.25
<b>The figure captions</b>	5	10	8	10	10	7	8	10	8.50	1.85
<b>Hypertext functionality</b>	8	10	10	10	10	9	10	10	9.63	0.74
<b>3d versus 2d presentation</b>	10	10	10	7	10	10	10	10	9.63	1.06
<b>The quality of the models</b>	3	4	5	3	3	5	4	3	3.75	0.89

Table 1: The figures given by the subjects in response to the 11 statements on the questionnaire.

The standard deviation figures (see Table 1) are relatively low in most cases, suggesting that the average values are a good reflection of the subjects' views. This includes the figure for the subjects' views on the quality of the models, which shows a very small standard deviation (0.89). The principal exception is the first scale (ability to view two illustrations at once) which shows a high standard deviation as well as the already-noted low average value. The result obtained in response to the statement regarding the integration of labels also shows a relatively high standard deviation, but the average figure itself is also high.

Since only one set of figures was obtained, it is not strictly appropriate to analyse these results using a standard test of significance. However, it is possible to gain some further insight into the nature of the results by comparing each score with the theoretical result which would have been obtained had there been no preference for either system (i.e., a score of five on each scale, the centre value). This was done using a Wilcoxon matched-pairs signed-ranks test. It was hypothesised that the scores would show a significant preference for or against each feature, but the direction of the preference was not predicted. Thus the results were analysed as for a two-tailed test. It was found that the subjects showed a significant preference ( $p \leq 0.02$ ) in favour of all the features tested except for two: the ability to view two models simultaneously, which showed no significant preference for or against; and the quality of the models, which showed a significant preference against.

The analysis of two users who agreed to be video-taped reveals that the newly developed interaction techniques can be used fluently and without help. It was interesting to see that medical students tend to enlarge the geometric model strongly to recognize even the smallest details. These scaling operations provide further evidence that there is not enough detail available.

## 4 Conclusions

The results show that the subjects regarded all the features evaluated here – with the exception of the simultaneous view facility and the quality of the models used – as useful to them in understanding the spatial relationships encoded in the geometric model. These features include the fisheye zoom to place labels and explanations in an illustration and the various ancillary features introduced in order to counter known or anticipated problems resulting from the use of fisheye techniques. The results therefore support the view that fisheye techniques can be used successfully, provided appropriate mechanisms are introduced to deal with the problem that nodes may disappear inadvertently. Moreover, the generation of figure captions and the automatic placement of labels, both distinguishing features of the Zoom Illustrator, were found to be useful in supporting the learning process.

The empirical evaluation helped us not only to assess the value of different interaction techniques but also to identify some problems and to assess them. The subjects, with their expert knowledge of anatomy, complained about the low level of detail available in the 3d-models used and wished for more information, e.g. about blood vessels and nerves, and more precise information about muscles and bones. More detailed geometric models are available, but these are an order of magnitude more complex than the ones used thus far and require very fast hardware for interactive handling.

A number of other problems were identified from subjects' comments. One is the lack of contrast between adjacent objects, in particular if both belong to the same organ-system and are therefore equally coloured. This is due to the rendering-methods used. To solve this problem, the edges of objects should be enhanced either by image-processing methods or by generating lines which indicate object boundaries. Concerning the labelling, it was sometimes said that it is not clear to which object a label refers or where an object referred to by a label actually starts and ends. The 1:1-connection between a label and an object is not always sufficient. Topographically complex objects, like muscles should be labelled with dedicated strategies to indicate their shape and branching structures. In [5] an approach to generate more appropriate annotations is described.

Moreover, it is interesting to note that students want to be able to underline parts of the texts and to make notes (e.g. hand-written remarks with a pen). They also wish more freedom to "play" with the geometric model, to remove parts and assemble objects themselves. We are thinking about designing a new mode for our learning system in which it can be used as a 3d-puzzle. Users can assemble objects themselves while the system supports them with docking points and snapping mechanisms to ease the 3d-interaction tasks involved.

Some of the features of the Zoom Illustrator deserve more thorough evaluation. In particular, the use of figure captions should be evaluated in detail with regard to the content and structure of the captions and the interaction facilities offered to influence the captions.

## 5 Acknowledgements

The authors wish to thank Marcel Goetze and Felix Ritter who conducted the tests and analysed the results. Furthermore, our thanks go to Thomas Strothotte for helpful comments on the design and evaluation of the Zoom Illustrator. Many thanks go to Christoph Thomas from the GMD, Sankt Augustin for his detailed comments on the experiment design and the statistical evaluation.

## 6 References

- [1] J. Sobotta (1988) Atlas der Anatomie des Menschen, 19. Edition, J. Staubesand (Ed.), Urban & Schwarzenberg, Munich-Vienna-Baltimore
- [2] A. Waldeyer & A. Mayet (1987) Anatomie des Menschen, Walter de Gruyter Verlag, Berlin
- [3] B. Preim, A. Raab & T. Strothotte (1997) "Coherent Zooming of Illustrations with 3D-Graphics and Textual Labels", Proc. of Graphics Interface, 19.-23. Mai 1997, pp. 105-113,
- [4] B. Preim (1998) Interaktive Illustrationen und Animationen zur Erklärung räumlicher Zusammenhänge, Dissertation, Otto-von-Guericke Universität Magdeburg, Fakultät für Informatik, VDI-Verlag, Reihe 10, Band 532
- [5] B. Preim & A. Raab (1998) "Annotation topographisch komplizierter 3D-Modelle", Proceedings of Simulation und Visualisierung, Magdeburg, 5-6th March 1998, SCS-Verlag, Delft, pp. 128-140
- [6] B. Preim, R. Michel, K. Hartmann & T. Strothotte (1998) "Figure Captions in Visual Interfaces", Proceedings of ACM Workshop on Advanced Visual Interfaces, L'Aquila, Italy, 24.-27. May), ACM Press, New York, pp. 235-246
- [7] G.W. Furnas (1986) "Generalized Fisheye Views", Proc. ACM SIGCHI'86 Conference on Human Factors in Computing Systems, Boston, April, pp. 16-23
- [8] J. Dill, L. Bartram, A. Ho & F. Henigmann (1994) "A Continuously Variable Zoom for Navigating Large Hierarchical Networks", Proc. of IEEE Conference on Systems, Man and Cybernetics, pp. 386-390
- [9] A. Raab & M. Rüger (1996) "3D-Zoom: Interactive Visualization of Structures and Relations in Complex Graphics", Proc. of Bildanalyse und -synthese, Erlangen, November, pp. 123-132
- [10] J.D. Mackinlay, G.G. Robertson & S.K. Card (1991) "The Perspective Wall: Detail and context smoothly integrated", Proc. of CHI '91, New Orleans, April, pp. 173-179
- [11] K. Hartmann, B. Preim & T. Strothotte (1998) "Describing Abstraction in Rendered Images through Figure Captions", Workshop on "Combining AI and Graphics for Intelligent User Interfaces of the Future", European Conference on Artificial Intelligence (ECAI '98), August, 1998, Brighton, UK (to appear)
- [12] K.-H. Höhne, B. Pflessner, A. Pommert, T. Schiemann, R. Schubert & U. Tiede (1996) "A Virtual Body Model for Surgical Education and Rehearsal", Computer – Innovative Technology for Professionals, January 1996, pp. 25-31.

## Authors' addresses

Dr. Ian Pitt  
 Department of Computer Science,  
 University College Cork,  
 Cork,  
 Ireland  
 Email: i.pitt@cs.ucc.ie

Dr. Bernhard Preim & Dipl.-Inf. Stefan Schlechtweg  
 Otto-von-Guericke Universität Magdeburg  
 Fakultät für Informatik  
 Institut für Simulation und Graphik  
 39016 Magdeburg  
 Email: {bernhard, stefans}@isg.cs.uni-magdeburg.de