# Perceptual and Cognitive Challenges to Learning with Dynamic Visualisations

Richard Lowe
Curtin University
GPO Box U1987 Perth WA
Australia
r.k.lowe@educ.curtin.edu.au

**Abstract:** Animated displays are becoming widely used as a means of teaching about dynamic subject matter. In recent years, advances in computer hardware and software have greatly facilitated the authoring of animations to the extent that their increasing use often appears to be driven more by technological than by pedagogical imperatives. As a consequence, there has been comparatively little focus upon issues concerning the requirements for effective instructional design of dynamic visualisations. Current development practices typically reflect a simplistic approach to the design of animations that has not progressed beyond the goal of providing dynamic realism. However, this approach is not necessarily consistent with the information processing characteristics of learners and their need for support in order to take advantage of animation's potential benefits. Principled approaches to the design of educationally effective dynamic visualizations are explored in the context of the complex domain of process control operation. Consideration is given to alternatives to realistic representations of the behaviour of complex dynamic systems that may be more consistent with learners' perceptual and cognitive capacities.

## Introduction

Much of the burgeoning educational use of dynamic visualizations appears to be based upon the conviction that these depictions are, by their very nature, more beneficial for learners than static visualizations. An interesting reflection of this attitude is that many of today's explanatory animations clearly have their origins directly in earlier static depictions of the same content. While there exists a widespread assumption that dynamic portrayals are intrinsically superior to their static counterparts, this view is highly simplisitic and not consistent with findings from recent research. Mayer (in press) emphasises the key role of instructional design in the use of such representations.

An unfortunate consequence of the prevailing preference for dynamic portrayals is the tendency amongst instructional designers to say 'we should animate that' rather than to ask 'should we animate that?'. Major drivers for this tendency seem to be both the increasing technical ease with which animated material can be generated and the pressure from market competition for 'attention-grabbing' visual effects in educational resources. Leaving aside the issue of the affective role that animations are supposed to play in learning, a common justification for using animations appears based on the 'directness' with which they can depict dynamic change. In other words, a good match between the representational medium and characteristics of the phenomenon being represented is considered instructionally desirable (e.g. Dijkstra, 1997). On the surface, this view appears very reasonable in terms of the comparative cognitive demands that dynamic and static visualisations are likely to make on the learner. In a general sense it can be argued that such matching should have information processing advantages for the learner because dynamic visualisations depict change processes explicitly rather than in the implicit fashion that is obligatory with static visualisations. Instead of having to perform the demanding process of generating a mental animation themselves (Hegarty, 1992), learners using dynamic visualisations should be able to 'read off' change information directly. By this argument, the use of dynamic rather than static visualisations should free up cognitive capacity for processing that is more central to the learning task.

## Matching and static graphics

The notion that instructional benefits will be obtained simply by providing a good match between the referent subject matter and its visual representation is of course not confined to animated depictions. It is also found with static graphics in cases where there is a preoccupation with depicting the subject matter in a realistic way. Until recently, photography was the standard way to produce realistic depictions of this type and for this reason such depictions were confined to concrete entities and real-world phenomena that could be captured photographically. However, the development of computer software capable of generating photorealistic images has meant that visualisers can now produce apparently realistic depictions of subject matter that cannot readily be photographed or that does not even have a physical reality.

A glance through a selection of present-day educational resources shows the current popularity of realistic treatments in static visualisation. As with animations, both affective considerations and technical advances are probably responsible for this current emphasis on realism. However, such an emphasis runs counter to the long established practice of simplifying pictorial representations intended for instructional or explanatory purposes. Simplification typically involves *reducing* the degree of realism in various ways in order to assist the learner. Approaches such as removing irrelevant material and making salient features of the subject matter more apparent were originally developed because the complexity of real-life situations often made it difficult to learn from them directly. By removing or changing some aspects of these situations via a graphic rendition, these approaches provided a better match between learners' limited information processing capacities and the presentation of the subject matter. The current fashion for matching the representation as closely as possible to its referent seems to run very much counter to such approaches and instead implies that the superficial characteristics of the referent situation (rather than the characteristics of learners) should primarily determine the characteristics of a depiction.

## Problems with animations

Realism in static visualization involves only the *appearance* of depicted content (visuospatial information). However, with dynamic visualization, there is the additional realism component of the *behaviour* of that which is depicted. Recent research indicates that dynamic depictions presenting realistic portrayals of the referents' behaviour may lead to their own problems for learners in terms of information processing (e.g. Lowe, 1999; 2001a; in press; Schnotz, Böckheler, & Grzondziel, 1999). Two distinct categories of such problems have been suggested to account for findings that show learners sometimes fail to benefit from animated presentations. One type of problem is 'underwhelming' in which animated presentation results in a reduction in the extent to which learners engage in valuable processing activities. The other and in some sense opposite type of problem is 'overwhelming' in which excessive information processing demands are imposed on the learner. With regard to this latter category, problems seems to arise because the animation is designed in such a way that its dynamics are dictated primarily by the exigencies of the situation being depicted rather than by the exigencies of the perceptual and cognitive processes necessary for effective learning. This predominant concern with verisimilitude can result in depictions in which 'realistic' presentation of aspects such as cause and effect, the speed of phenomena, and layered interrelations of functions are so demanding that they impair the learner's capacity to process the available information effectively. This type of impairment would seem to be particularly likely for animations that present complex subject matter but make little or no concession to the capacities of the target learners.

It is clear that access to sufficient information about dynamic characteristics of the subject matter is absolutely vital for building satisfactory mental models of dynamic situations. However, this does not necessarily mean that such temporal information should be presented as realistically as possible (any more than visuospatial information should be presented realistically in static displays). The core requirement is to provide the relevant information about situational dynamics in a manner that is highly accessible to the learner; this does not require that it is presented in the most direct analog form possible. As mentioned earlier, static depictions of dynamic situations can have their own drawbacks as tools for learning because they require learners themselves to extract salient dynamic cues and mentally animate the presented information. This can be a demanding process and one that is subject to errors when complexities such as extended cause-effect chains are involved (see Narayanan & Hegarty, 2000). However, the solution to the problems that can occur with static displays is not simply to turn them into animations that provide realistic depictions of the behaviours involved. Nevertheless, this approach seems to have been adopted by many of those who are currently producing animations from their static precursors.

## Realistic appearance or behaviour?

An interesting feature of the problem of 'overwhelming' described above is that it seems to occur even with dynamic material that is realistic only in terms of the behaviour it depicts. In other words, static depictions that have been greatly simplified in a visuospatial sense may nevertheless be overwhelming if their dynamic characteristics are presented in a highly realistic manner when they are converted into animations. The animated weather map displays used in the research that revealed this problem (Lowe, 1999) are largely comprised of mathematically-derived abstract graphical representations that cannot be matched in any direct way to what we 'see' of the weather in our everyday surroundings. Their depictions of information about the weather are far from realistic in terms of appearance. In a visuospatial sense, weather maps present a highly simplified version of the referent system that has considerable advantages for those whose job it is to interpret meteorological information. Animated versions of the static weather maps usually employed to teach meteorology were developed with the intention of facilitating students'construction of dynamic mental models (Lowe, 1995). However, it was found that even when visual cues and a high level of user control were provided, learners were relatively unsuccessful in extracting meteorologically relevant information from the animation. These findings raise the possibility that some of the animations used in educational materials are far less pedagogically effective than their originators may assume. Even when these animations are produced by graphics professionals, their design is unlikely to be based upon a principled understanding of what is required to learn from dynamic visualisations. While many of these artists undoubtedly have advanced animation skills, the particular nature of their skills may not always be very well suited to the needs of education. The extensive body of existing animation knowledge that is used to train today's graphic designers has largely been built up from the work of animators in the entertainment industry whose objectives and measures of success are somewhat different from those of educators. Consequently, the tried and true approaches developed for entertaining us with the animated exploits of Mickey Mouse or Felix the Cat will not necessarily be those that are particularly effective for facilitating the learning of academic content.

While the large bulk of instructional animations produced today are the result of attempts to provide more explicit educational explanations, this is not always the case. In a variety of subject domains, the material facing the learner is by its nature already dynamic and the learner's task is to respond to the depicted changes appropriately. This is often the case with industrial training, particularly in situations where the instructional goal is to develop sophisticated real-time operational capacities in the trainee. In a typical situation of this type, the trainee is confronted by some kind of continually changing visual display and must learn to read that display, interpret the dynamic situation it depicts, then take the necessary action. An example of this type of instruction occurs when trainee operators are being taught to use the process control displays that are widely used in extractive and manufacturing industries for the remote monitoring and control of production plant. These computer-based displays are dynamic diagrammatic representations that have been designed to inform experienced operators about the operation of the plant and allow them to adjust its component processes. The design features of these displays are intended to maximise the effectiveness of the seasoned operator, not to teach novices the skills required to perform these operational functions. In this sense, a process control display differs markedly from the more conventional types of animated display used in learning materials. However, its distinctive characteristics provide a valuable opportunity for considering educational aspects of animations from a new and revealing perspective. Of particular interest is the challenge of providing a sufficiently 'authentic' form of instruction for process control trainees in a task that is very demanding in terms of its perceptual and cognitive requirements. In the next section, a recently developed set of training materials will be described and used to raise some key issues about the instructional design of dynamic visualisations.

## Dynamic Visualisation in Process Control Operation

The *Smart Operator*® package provides training in the operation of process control systems and is based upon accurate real-time simulations underpinned by an expert system that monitors trainee performance and provides formative feedback (Lowe & Neilsen, 1998; Lowe, 2001b). Figure 1 is a screen shot from *Smart Operator*® showing the way information about a plant's operation is represented on the display. However, in contrast to this static illustration, the actual display depicts information in a dynamic fashion. There are continual changes in the material shown on the screen and an important goal of *Smart Operator*® is to teach trainees how to deal with this ever-changing set of visual information. The simulation presents a display that, in terms of both appearance and behaviour, is a close approximation to an actual process control screen. It should be noted however that, in contrast to the *behaviour* of the display, the *appearance* of the material on a process control screen is highly *unrealistic* with respect to the situation it represents because it is a greatly simplified depiction of its referents. The abstract graphic entities shown on these screens represent the equipment comprising the manufacturing plant that is to be controlled by the process operator. In the plant itself, individual pieces of equipment are fitted with electronic sensing and control devices that allow their operations to be monitored and adjusted at a distance. These sensors and controllers are connected by electric cables to the plant's central control room where they are linked to a computer display. This process control display (a) provides information about the operational state of the plant's equipment and (b) allows the process control operator to make adjustments to individual items of that equipment. Various components of the display provide the operator with dynamic information about changing aspects such as temperatures, pressures, levels and flow rates. Even for a relatively small plant, these displays are very complex and require highly skilled operators. It is crucial that these operators are able to read, monitor, interpret then appropriately respond to the displayed graphics in order ensure the plant operates safely and effectively.

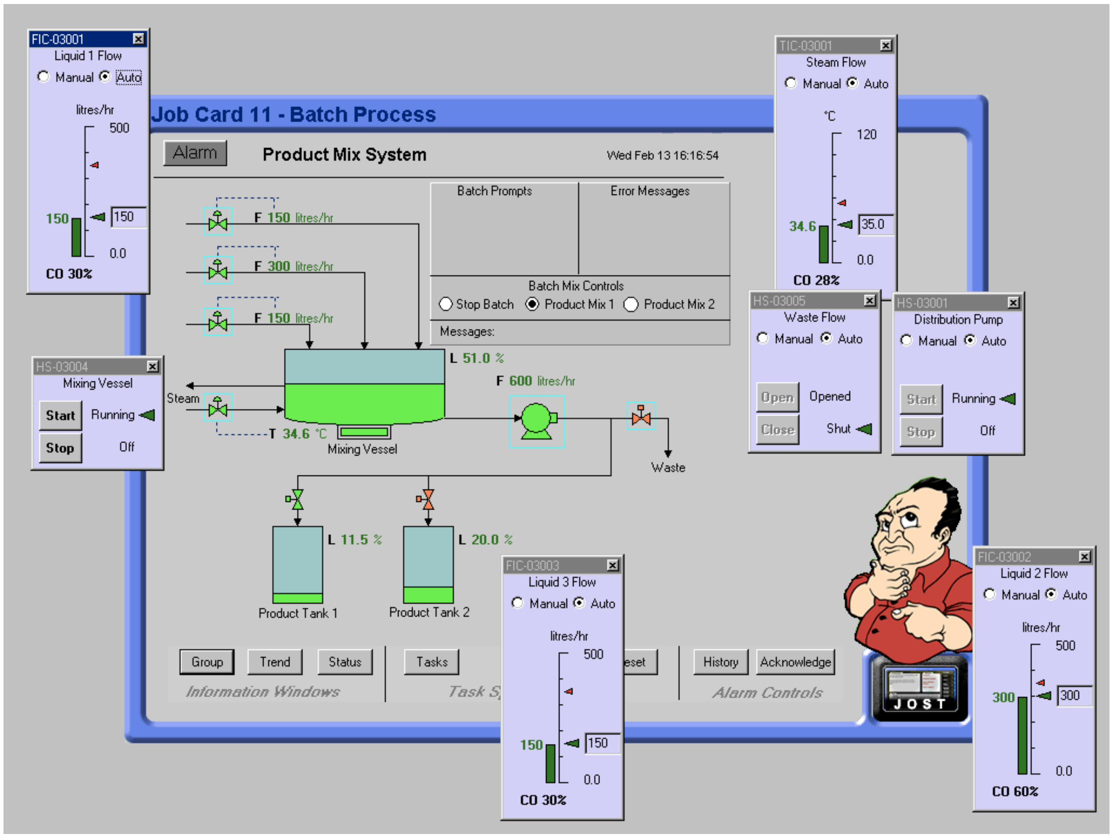


**Figure 1.** A typical *Smart Operator®* display with diagrammatic representations of equipment and faceplates.

Presenting the fundamentals of process control operation to trainees in a readily accessible way is a very challenging instructional design task. In addition to the inherent complexity of the systems represented, the on-screen information pertaining to those systems is presented in a variety of ways (symbolic, graphical, textual, numerical, etc.). Relevant aspects of this diverse information must be extracted and synthesised in order to interpret the overall situation correctly. Real-life process control displays present their information as economically as possible and make no allowance for the interpretative difficulties that this parsimony can pose for those without experience of these depictions. Important information provided by a process control screen is often hidden (accessible only by calling up additional displays such as faceplates) or visually inconspicuous. These types of features present considerable problems for the process control trainees who may not even be able to find vital information signals generated by the display, let alone make all the connections between items of information necessary to respond appropriately. Key instructional design challenges are to help learners (a) search the displays productively, (b) recognise those aspects of displays that present crucial information, and (c) relate the extracted information in ways that make the depicted situation comprehensible.

## Instructional design considerations

A major constraint on designing effective training in the fundamentals of process control is the need to balance (a) the desirability of having 'near' rather than 'far' transfer from the training materials to actual process control displays with (b) the challenge of helping learners handle the high level information processing demands of realistically-presented process control displays. The design must therefore be a trade-off between providing a training environment that bears a sufficiently close resemblance to a real process control situation and introducing sufficient modification of the authentic situation to facilitate learning. Industry trainers agree that novices find actual process control screens overwhelming and yet are adamant that one of the first things trainees need to learn is how to cope with this complexity. It is therefore necessary to consider carefully the nature of the task that faces learners and determine the specific information processing challenges involved. To this stage, there has been a lack of formal empirical investigation of the basic demands of reading process control displays. However, these displays share a number of fundamental characteristics with weather maps such as their conceptual complexity, abstraction, and role as tools for making predictions about remote situations. For this reason, it is useful to explore what implications findings from research into the processing of weather map diagrams may have for designing training in the interpretation of process control displays. The following discussion begins by considering investigations based upon static weather maps and then moves on to consider weather map animations.

An important foundation for working with process control screens is the capacity to subdivide the presented material appropriately according to the different classes of equipment represented and operational information concerning that equipment. These screens typically depict several interconnected process loops that represent the particular functions carried out by different types of equipment (pressure loops, level loops, temperature loops, flow loops, etc.). Understanding process control operations relies upon being able to identify and distinguish between these different subsets of information and this requires the viewer to segment the display into appropriate functional chunks. In the field of meteorology, this type of chunking task is one that is done poorly by those who lack sufficient domain-specific background knowledge (Lowe, 1993; 1996). Subjects without specialised meteorological training or experience tended to group weather map markings on the basis of their visuospatial characteristics rather than on the basis of domain-related criteria. In contrast to professional meteorologists, these novices grouped markings according to characteristics such as figural similarity and proximity instead of their meteorological identity and relationships. A possible implication for designing process control training is that learners would require explicit direction as to how process control displays should be subdivided.

Effective use of a process control display requires the operator to switch between global and local aspects of screens as different types of operational need arise. Research on novices' processing of static weather maps suggests that their lack of meteorological background knowledge prevents them from dealing with the presented information in the sophisticated multi-layed manner that is typically of professional meteorologists (Lowe, 1994). In addition to their superficial approach overall, novices tended to deal with localised fragments of information individually rather than connecting highly related yet widely separated components of the display. Their preoccupation with local structures amongst graphic elements was accompanied by a corresponding neglect of structures existing on a more global scale. In process control, it is essential for operators to synthesise items of information that are not only scattered across the screen but are also presented via very different formats (symbols, graphs, text, numbers, etc.). In addition to giving trainees explicit guidance in how to subdivide process control displays, it may also be necessary to provide cues that help them to link together items that, due to their wide separation and superficial differences, appear not to be related to each other (see also Winn's 1993 discussion of the configuration in diagram search).

The information-rich, dynamic character of a process control screen presents trainees with a situation in which selective attention is needed in order to 'read' the screen effectively. This selectivity is required across both space and time. Investigation of how novices deal with static weather maps indicates that their attention tends to be captured by features that are graphically distinctive while meteorologically important aspects that lack this distinctiveness remain unnoticed. As is the case with weather maps, there is no simple correspondence between how conspicuous graphic elements are in a process control display and their importance. The task of learning to find relevant information on a process control screen is further complicated by the dynamic nature of these displays. Although changes in the display are sometimes deliberately used to alert the operator to important shifts in system status (such as the flashing red alarms that indicate a problem condition), the majority of screen changes are simply a reflection of the on-going processes. The significance of these changes generally bears no relationship to their graphic distinctiveness. For the trainee, there are virtually no visual cues about what to attend to and what to ignore. Work with animated weather maps (Lowe, 1999; 2001a; in press) suggests that even when novices are given extensive control over factors such as the playback speed, direction and frequency of animations, they are relatively poor at isolating temporal fragments that contain task-relevant information. In process control operation, the flux of information presented on the display is largely driven by changes occurring in the plant's equipment and so the operator has comparatively little opportunity to control its presentation as was the case with the weather map animations, As a result, it could be expected that novices' capacity to extract relevant information from a process control screen presenting the behaviour of the system in a realistic way would be very limited.

## Discussion

The process control example described in this paper demonstrates that tension can exist between the desirability of providing dynamic displays that capture the behaviour of a system in an authentic manner and the demands such dynamic authenticity make on the perceptual and cognitive capacities of learners. The challenge for instructional designers who must devise introductory learning experiences for such systems is to balance these competing demands in a way that efficiently develops learners' abilities to deal with the dynamic complexity in an effective manner. The standard visuospatial approaches used with static instructional illustrations of removing 'inessential' material from a depiction and adding visual cues to direct attention to salient information appear to have limited applicability in this context. Negative effects that may result include disruption of the overall appearance of the display, removal of potentially important relationships, and imposition of a further graphic layer to complicate the interpretation process. Research on weather map animations suggests that while the addition of visual cues can have positive effects in terms of directing learners' attention to salient parts of the display, it may also have unforeseen negative consequences with regard to the way they interpret relationships between dynamic entities (Lowe, 1999).

A design approach that may hold more promise as a way of limiting the perceptual and cognitive demands made by such animated displays is to give more emphasis to temporal manipulation of the display than to visuospatial manipulation. It has already been noted in the earlier consideration of weather map animations that simply providing high degree of user control over the temporal dimension does not seem to be particularly effective in facilitating novices' extraction of salient information from a dynamic display. A subsequent study (Lowe, 2001a) indicates that the reason may be that such learners who lack background knowledge in a domain do not interrogate the temporally distributed information in an effective manner. Although occurring in a different dimension, this deficiency in learners' search of a dynamic visualisation is somewhat analogous to the deficiency found in the way they search for visuospatial information in a static display. In the case of static displays, it appears that the difficulties learners have in extracting information required for the proper interpretation of the depiction can arise from factors such as their inability to identify informationally-important regions of the display. This identification process often relies on background knowledge that novices simply do not possess. There may be a similar explanation for problems that learners have in identifying potentially informative *temporal* 'regions' within an animated sequence. In this dynamic case, these problems may be compounded because of the transitory nature of the images presented and the presence of dynamic perceptual cues that are additional to the visuospatial cues present in static depictions. Graphic entities in animations that contrast greatly with their dynamic context are not necessarily those that are the most thematically relevant. Simply giving learners a high degree of control over the playback of dynamic visualizations is not therefore necessarily likely to lead to productive interrogation and interpretation of those displays.

Because learners who are novices in a domain are poorly equipped with the background knowledge required to undertake strategic temporal interrogation of domain-related animations, it is likely that instructional designers need to provide them with extensive support to guide their interrogation in productive directions. This type of support will be described as *temporal cueing*. It could include approaches such as the prior subdivision of animations into separate temporal segments that isolate salient aspects of the events depicted or the freezing of the action at certain strategic points in the animation that contain information of high thematic relevance. Further possibilities are the imposition of appropriate changes in pace and direction that help to reveal small or large scale temporal patterns that novices may not attend to if the behaviour of the dynamic situation is represented as it actually occurs. Such strategies are likely to be particularly important in developing interpretation of cause-effect relationships. The aim of these manipulations should be to address the perceptual and cognitive challenges that can be associated with dynamic visualizations in ways that as far as possible preserve the inherent character of the display. Excessive distortion of the dynamic behaviour may engender the type of interpretative problems that were found to occur when intrusive forms of visuospatial cueing were applied to animations. Research is clearly needed to establish the relative effectiveness of different types or levels of temporal cueing and the way that such cues may interact with visuospatial forms of cueing.

## Acknowledgment

My sincere thanks to CBT Simlulations Technology for the opportunity to work with *Smart Operator*®.

## References


Dijkstra, S. (1997). Educational Media and Technology. In S. Dijkstra, N. Seel, F. Schott & R. Tennyson (Eds.), *Instructional design: international perspectives* (Vol. 1, pp. 137-143). Hillsdale, NJ: Lawrence Erlbaum Associates.

Hegarty, M. (1992). Mental animation: inferring motion from static diagrams of mechanical systems. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *18*, 1084-1102.

Lowe, R.K. (1993). Constructing a mental representation from an abstract technical diagram. *Learning and Instruction*, *3*, 157-179.

Lowe, R.K. (1994). Selectivity in diagrams: reading beyond the lines. *Educational Psychology*, *14*, 467-491.

Lowe, R.K. (1995, July). Developing basic chart reading skills: using interactive animation for building effective mental models. Paper presented at the *2éme Conférence Internationale sur l'Enseignement Assisté par Ordinateur et l'Enseignement á Distance en Météorologie,* Tolouse, France.

Lowe, R.K. (1996). Background knowledge and the construction of a situational representation from a diagram. *European Journal of Psychology of Education*, *11*, 377-397.

Lowe, R.K. (1999). Extracting information from an animation during complex visual learning. *European Journal of Psychology of Education*, *14*, 225-244.

Lowe, R.K. (2001a, August). Learners' search processes during exploration of interactive animation. Paper presented at the *9th European Conference for Research on Learning and Instruction*, Fribourg, Switzerland.

Lowe, R.K. (2001b). Supporting effective reading of pictorial materials in visually oriented learning environments. In F. Lockwood and A. Gooley (Eds.). *Innovation in Open and Distance Learning* (pp. 201-212). London: Kogan Page

Lowe, R.K. (in press). Animation and learning: selective processing of information in dynamic graphics. *Learning and Instruction*.

Lowe, R.K., & Neilsen, T.E. (1998, June). Design and development of an intelligent simulation training system for process control operators. *Proceedings of Ed-Media/Ed-Telecom 98 Conference* (pp. 839-844), Freiburg: Association for the Advancement of Computing in Education.

Mayer, R. E. (in press). The promise of multimedia learning: using the same instructional design methods across different media. *Learning and Instruction.*

Narayanan, N.H., & Hegarty, M. (2000). Communicating dynamic behaviors: are interactive multimedia presentations better than static mixed-mode presentations? In M. Anderson, P. Cheng & V. Haarslev (Eds.). *Theory and Application of Diagrams* (pp 178-193). Heidelberg: Springer.

Schnotz, W., Böckheler, J., & Grzondziel, H. (1999). Individual and co-operative learning with interactive animated pictures. *European Journal of Psychology of Education*, *14*, 245-265.

Winn, W.D. (1993). An account of how readers search for information in diagrams. *Contemporary Educational Psychology*, *18*, 162-185.


.

.