Problem Effectiveness in a Course Using Problem-based Learning

DIANA H. J. M. DOLMANS, MA, WIM H. GIJSELAERS, PhD, HENK G. SCHMIDT, PhD, and SYB B. VAN DER MEER, MD

Background. Problem-based learning (PBL) emphasizes active generation of learning issues by students. Both students and teachers, however, tend to worry that not all important knowledge will be acquired. To explore this question, problem effectiveness (i.e., for each problem, the degree of correspondence between student-generated learning issues and preset faculty objectives) was examined in three interdependent studies. **Method.** The three studies used the same participants: about 120 second-year students and 12 faculty tutors in a six-week course on normal pregnancy, delivery, and child development at the medical school of the University of Limburg in The Netherlands, 1990–91. The participants were randomly assigned to 12 tutorial groups that were each given the same 12 problems; the problems were based on 51 faculty objectives; the tutors were asked to record all learning issues generated by their groups. Study 1 addressed this question: To what degree are faculty objectives reflected by student-generated learning issues? Study 2: To what extent do students miss certain objectives, and are these objectives classifiable? Study 3: Do students generate learning issues not expected by the faculty, and are these issues relevant to course content, and finally, why do students generate these issues? To help answer these questions, the studies employed expert raters and a teacher familiar with the course content. **Results.** Study 1: For the set of 12 problems, the average overlap between learning issues and faculty objectives was 64.2%, with the percentages for individual problems ranging from 27.7% to 100%. Study 2: Of the 51 objectives, 30 were not identified by at least one tutorial group; these objectives were grouped into three categories; on average, each group failed to identify 7.4 objectives (15%). Study 3: Of 520 learning issues, 32 (6%) were unexpected; 15 of these were judged to be at least fairly relevant to course content; they were grouped into four categories. **Conclusions.** The students' learning activities covered an average of 64% of the intended course content; in addition, the students generated learning issues not expected by the faculty, and half of these issues were judged relevant to the course content. Thus, PBL seems to permit students to adapt learning activities to their own needs and interests. *Acad. Med.* 68(1993):207–213.

The concept of problem-based learning (PBL) is generally seen as an attempt to apply modern cognitive views on teaching and learning to instructional practice and curriculum design.[1] Many investigations indicate that competence is fostered not primarily by teaching to transmit information but through teaching to encourage specific kinds of cognitive

activities.[2] Learning needs to be an active, constructive mental activity that enables students to build from what they already know. Students should become architects of their own knowledge and eventually take full responsibility for their own learning. Problem-based curricula are assumed to reflect these scientific conceptions of the conditions of human cognition and learning.[1]

First, it is assumed that PBL promotes contextual learning, since new knowledge is acquired in the context of some meaningful problem or situation. Situated knowledge is considered to be more accessible because the situational cues that activate the knowledge are stored within the same cognitive structures.[3,4] Second, discussion of a problem will activate previously acquired knowledge in order to construct a mental representation or model of the problem, which, in turn, would facilitate the

comprehension of new information relevant to that particular problem.[5,6] Third, problem discussion will increase intrinsic motivation because it involves the learners more actively in the issues at hand.[5] And fourth, the emphasis on self-directed learning will encourage students to reflect upon and control their own learning activities and develop self-regulatory skills conducive to lifelong learning.[2,5]

In PBL students are confronted with a problem consisting of a description of a set of phenomena in need of explanation. The problem has been specifically designed for instructional purposes. Eight to ten students, in a group guided by a tutor, discuss the problem and try to explain the phenomena in terms of their underlying processes, principles, or mechanisms.[7] During this initial discussion, issues emerge that require further exploration. These student-generated learning issues are topics

that the individual tutorial group decides are prerequisites for a better understanding of the problem under discussion.[8,9] Thus, learning issues are the starting points for students' learning activities and serve as guides for studying literature relevant to the problem at hand. Students engage in self-directed learning activities in order to find information regarding the issues raised. In a second tutorial session, each participant reports the information he or she has found. During this meeting an attempt is made to integrate the information that has been collected and to draw conclusions.

When teachers design problems, they have in mind certain topics that students are expected to cover. Hence, a goal of curriculum design in this area should be to construct problems that are indeed effective in reaching the intended learning outcomes. Whether students undertake learning activities planned by faculty members is to a large extent determined by the learning issues generated. If students fail to generate the appropriate issues, the preset objectives are not identified, and hence the intended learning outcomes are not accomplished. Ineffective problems cause difficulties for students in generating the appropriate learning issues, implying that students' learning activities may not cover the intended subject matter, which, in turn, will lead to lack of content coverage. Consequently, the effectiveness of a problem may be defined as the degree of correspondence between student-generated learning issues and preset faculty objectives. This definition is illustrated by a problem presented to the students at the medical school of the University of Limburg, The Netherlands, during a second-year course about normal child development in the 1990–1991 academic year.

Problem: Tall Girl

During the last few years, Ellen has grown tall very quickly. She has always been a tall girl, but at the age of 11 years and the height of 5 feet, 4 inches, she rises head and

shoulders above her age group. People always take her to be older, which sometimes becomes wearisome. What will become of her? She still has not reached puberty

In designing this problem, the teachers had five objectives in mind: (1) what are the normal rates of child growth, (2) what are the normal stages in secondary sexual characteristics, (3) what endocrine control processes influence growth, (4) what are the psychological effects of being extremely tall compared with others in one's age group, and (5) what diagnostic procedures are available to predict ultimate height. Given the problem, one tutorial group generated five learning issues: (1) what are the normal patterns of growth rate, (2) what factors influence growth, what hormones are concerned with growth, and what are their effects, (3) what kinds of physical changes take place during puberty and what are the possible explanations, (4) what tests can be used with regard to growth, and (5) what medical interventions are possible in a case of abnormal growth.

Comparing the learning issues of this tutorial group with the preset faculty objectives in detail reveals that this tutorial group only partially covered the faculty objectives. Four of five objectives were identified by the tutorial group. The fourth objective, addressing the psychological effects of being extremely tall, was not identified by this tutorial group. This is a Type A mismatch. Interestingly, the tutorial group generated one learning issue, what a doctor can do with a case of abnormal growth (issue 5), that did not correspond to any of the faculty objectives and as such was not expected. This is a Type B mismatch.

Whereas problem effectiveness can be expected to determine students' learning, only a few studies have actually addressed the relationship between the nature of problems and students' learning activities.[10–12] The purposes of the present article are (1) to provide an overview of previous research conducted in this area and (2) to describe three interdependent studies of the overlap between faculty

objectives and student-generated learning issues.

REVIEW OF THE LITERATURE

Consideration of the effectiveness of problems in PBL is dominated by the question whether students are able to generate all the appropriate learning issues and are able to identify the required knowledge. Faculty members' observations with regard to the effectiveness of problems suggest that students sometimes tend to focus on just one detail of the problem that is most appealing to them.

The PBL literature contains several examples of studies that examine whether students actually identify the knowledge required.[10–12] These studies have investigated the matches between course objectives and student-generated learning issues. Tans et al.,[12] for example, studying a problem-based course on muscle physiology for physiotherapy students, used two techniques to analyze the match between faculty objectives and student's learning issues: (1) a method based on experts' judgements and (2) a text-reduction method. The first method consisted of requiring experts to judge for each faculty objective whether it was reflected in learning issues. This method resulted in an average percentage of overlap of 68.3 (SD, 24.6). The second method, a text-reduction technique, consisted of reducing both faculty objectives and learning issues into lists of concepts and comparing the lists to determine the overlap. The percentage of overlap between the two lists of concepts was 39.6 (SD, 22.9). The authors argued that the text-reduction technique is more stringent in determining the percentage of overlap than is the technique of using experts' judgements. The difference between the two techniques may be explained by the fact that text reduction is based on the assumption that two concepts correspond with each other only if exactly the same terminology is used, whereas experts tend to interpret what is intended rather than to focus on precise formulations.

Coulson and Osborne[10] also con-

ducted a study to examine to what extent students were able to identify all course objectives. With specific objectives in mind, a couple of teachers selected seven patient-problem simulations in a PBL module. Learning issues generated by the tutorial groups were collected and three raters compared them with a faculty-generated list of objectives. The results showed that, for all the tutorial groups taken as a whole, the groups succeeded in identifying all course objectives. However, at the individual group level, each faculty objective was identified on the average by five of the twelve tutorial groups (42%). No information was provided as to whether the unidentified objectives were important.

Students' success in identifying faculty objectives was also tested in a study conducted by Shahabudin.[11] In this study, a conventional course was compared with a parallel problem-based track. The former course consisted of lectures in systemic pathology and clinical microbiology as part of a conventional medical curriculum. The problem-based course included a ten-week period during which tutorial groups discussed a series of problems selected from the conventional systemic pathology course. Each tutorial group was asked to list all the learning issues that arose from the discussion. At the end of the period, all learning issues were compiled into a smaller set of more general subjects. This list of subjects was compared with the faculty members' list of topics to be covered by the tutorial groups and the list of topics to be covered in the conventional curriculum. It was found that, in general, the tutorial groups succeeded in identifying all required topics. The author concludes that this congruence indicates that learning from problems is not entirely student-directed but can largely be influenced by teachers.

The results of the studies discussed thus far indicate that students succeed reasonably well in identifying preset faculty objectives. The studies fail, however, to suggest why some faculty objectives are not always identified by students. Exploring the nature of these unidentified objectives, or searching for their common characteristics, can probably give more insight into why they were not identified. Furthermore, the studies discussed do not provide any hint as to the nature of the unintended learning issues. Tans et al.[12] conclude that a majority of the unexpected learning issues were relevant to the course content, without further specifying why these issues were generated.

THREE STUDIES

The purpose of the present investigation was to analyze the relationship between preset faculty objectives and student-generated learning issues, both with regard to faculty objectives not identified by students (Type A mismatches) and with regard to learning issues not expected by teachers (Type B mismatches). Three studies were designed to suggest answers to three sets of questions: (1) To what degree are faculty objectives reflected by student-generated learning issues? (2) To what extent do students miss certain objectives, and are these faculty objectives classifiable? (3) Do students generate learning issues not expected by the faculty? Are these issues relevant to the course content, and why do students generate these issues?

All three studies used the same group of participants: about 120 second-year students at the medical school of the University of Limburg in The Netherlands, 1990–1991. The medical school has a problem-based curriculum. In addition, 12 faculty tutors were part of the experiments. Both the students and the tutors were randomly assigned to 12 tutorial groups. All analyses were conducted at the level of the tutorial group.

The three studies also used the same basic data. The course under study was a six-week course that covered themes related to normal pregnancy, delivery, and child development, represented by 12 problems about topics such as childbirth, psychomotor development, and psychosexual development. The problems were based upon 51 faculty objectives that covered the course as a whole; for individual problems the numbers of objectives ranged from one to eight, with an average of 4.3. For each problem, the tutors were asked to record all learning issues generated by their tutorial groups. The tutorial groups generated one-to-ten learning issues per problem, with an average of 3.6.

STUDY 1

Study 1 was designed to assess to what degree preset faculty objectives were reflected by student-generated learning issues. This study can be considered a conceptual replication of the Tans et al.[12] study. Experts' judgements were used to measure the concordance between course objectives and student-generated issues.

Procedure

Twelve pairs of 24 expert raters were asked to judge the matches between faculty objectives and student-generated learning issues. For a given problem and a given tutorial group, the pairs were instructed to compare each student issue with each objective, with each rater individually judging the correspondence between faculty objectives and learning issues. For example, a problem might have contained five objectives, and the tutorial group might have generated five learning issues. Hence, each rater in the pair would have had to make 25 comparisons for a single problem and a single group. To reduce the total number of comparisons to be made by each pair of raters, a design was used in which the raters were nested within problems and within tutorial groups, i.e., each pair rated each group's performance on only one problem. (If each pair has been required to judge all learning issues generated for all 12 problems by each of the 12 tutorial groups, each individual rater would have had to make about 2,200 comparisons).

For each comparison, each rater had to judge whether a particular learning issue (1) definitely corresponded with a particular faculty objective or (2) definitely did not corre-

spond with that objective. (For the purpose of analysis, the few cases when raters could not give definite judgments were considered noncorrespondences). A faculty objective was assumed to have been identified by the students when the two raters agreed that a learning issue corresponded with the objective. If both raters agreed that there was no correspondence or if the raters disagreed (i.e., one indicated a correspondence and the other a noncorrespondence), then an objective was assumed not to have been identified by the students (presumably, they had failed to study the subject matter specified by the objective).

Results

The levels of interrater agreement for the individual pairs of raters ranged from 67.5% to 97.3%, with an average of 78.9%. Since these percentages might have been overestimated true interrater agreement, because of agreements based on chance alone, Kappa coefficients were computed for each pair of raters as an alternative measure for interrater reliability. The Kappa coefficient takes into account the expected proportion of agreement due to chance alone. It can take values between zero and one.[13] The Kappa coefficients for the pairs of observers ranged from .23 to .64, with an average of .43. The average Kappa value of .45 demonstrates a moderate level of agreement between raters.

It should be noted that the correspondence between faculty objectives and student-generated issues is quite difficult to judge, since a learning issue arises from a discussion in the tutorial group that, in itself, is not available to the raters. This lack of background information in some cases makes it difficult to judge to what extent a learning issue is relevant to a particular objective, although the students involved probably know exactly what is intended by the learning issue. The level of agreement between the raters, however, was fairly high, indicating that in most cases the raters were able to make clear judgements, even without additional context information.

Given the results of both methods for assessing reliability, the conclusion may be drawn that the rating procedure provides fairly reliable ratings.

The effectiveness of the problems was measured as the degree of overlap between faculty objectives and learning issues. As shown in Figure 1, the degrees of overlap between objectives and issues for the 12 problems ranged from 27.7% to 100% ($F_{(11,143)} = 6.84$, $p < .000$). The average overlap for the set of 12 problems was 64.2%. For ten of the 12 problems, the overlaps were close to this average (within 15 percentage points).

In summary, faculty objectives identified by the 12 tutorial groups averaged about 64%. However, the percentages differed across problems, implying that some problems were more effective than others.

STUDY 2

The purpose of Study 2 was to explore why some tutorial groups did not identify certain faculty objectives. Thus, the attention in this study was focused on those faculty objectives definitely not identified by one or more tutorial groups (i.e., Type A mismatches).

Procedure

To investigate why some tutorial groups had failed to identify certain faculty objectives, further analyses were conducted. Faculty objectives not identified by one or more tutorial groups were ranked in a list, based on the numbers of tutorial groups omitting them. This list was presented to a teacher who was familiar with the course content (he was the teacher responsible for the design of the course). He was asked to categorize the objectives listed, using any classification he considered meaningful.

Results

The results of the match procedure described in Study 1 showed that 30 of 51 faculty objectives were not identified by one or more tutorial groups. The numbers of tutorial groups failing to identify one or more of these
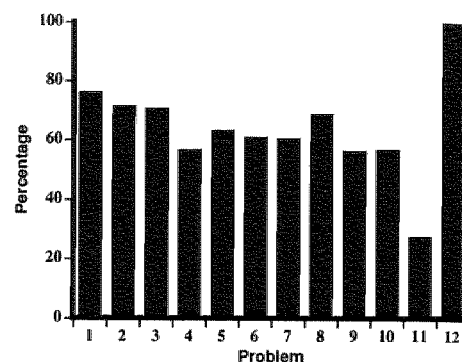
Figure 1. **Percentages of overlap between preset faculty objectives and student-generated learning issues for problems in a second-year course on normal pregnancy, delivery, and child development, University of Limburg, The Netherlands, 1990–91. About 120 students, assigned to 12 tutorial groups, participated in the course. The overlap between learning issues and faculty objectives was determined by pairs of expert raters.**

objectives ranged form one to eight. Nine of the faculty objectives were neglected by one tutorial group; another nine objectives were missed by two tutorial groups. Only two objectives were neglected by eight of the 12 tutorial groups. On average, each tutorial group failed to identify 7.4 objectives. This implies that about 15% of the course objectives (7.4 of 51) intended by the teachers were definitely not identified by the students as relevant to their learning.

The teacher's classifying procedure revealed three major categories of Type A mismatches: (1) faculty objectives related not so much to problems as to other curricular activities, (2) objectives spanning more than one problem, and (3) psychological and social objectives. Although these categories are not mutually exclusive and exhaustive, they suggest why these faculty objectives were not identified by the students.

The first category consists of faculty objectives to be attained through curricular activities other than small-group work. It should be noted that the medical school at the University of Limburg not only tries to stimulate learning through problems, but also offers opportunities to master the

medical domain through other instructional means. Some curricular objectives are to be achieved through laboratory work. Others are mastered through skills training such as physical examination training. Students participating in these curricular activities might have found it unnecessary to generate learning issues with regard to these objectives, because they had already studied the subject matter in the context of other curricular activities. For example, training was offered to the students in which they learned how to examine the position of the fetus. One problem presented to the students also dealt with prenatal examination, and the students were intended to study the different positions of the fetus. It appeared that some students did not identify this objective.

The second category consists of objectives broadly defined and requiring an extensive search into the literature. Examples of such objectives are "the influence of the environment on child development" and "theories of child development." These objectives hardly guide students in studying a specific topic, since they refer to very broad issues.

The third category consists of psychological and social faculty objectives, such as "psychological effects of being extremely tall compared with the age group" or "psychological well-being of a woman during pregnancy." The first example reflects an objective of the tall girl problem described earlier. This objective was not identified by eight of the 12 tutorial groups.

Furthermore, the data suggest that, if faculty objectives of a certain problem focus on both biology and psychological objectives, students generate biology learning issues at the expense of psychological issues. By contrast, if a problem contains only psychological objectives, students do generate these issues and do identify the required course objectives.

STUDY 3

Self-directed learning implies that students are free to choose their own learning issues. As a consequence, they may generate issues not ex-

pected by the faculty. These learning issues may still fit within the framework of the course, since PBL accords students an active role in the acquisition of knowledge. Study 3, therefore, concentrated on the nature of the issues generated beyond faculty objectives (i.e., Type B mismatches). The guiding questions of study 3 were (1) To what degree do students generate unexpected learning issues? (2) To what extent may these issues be considered relevant to the course as a whole? (3) Why do students generate these issues?

This study highlights student-generated learning issues rather than faculty objectives. The data consisted of the student issues judged as "not reflecting any of the faculty objectives." If both raters agreed that a student-generated learning issue definitely did not correspond to any of the faculty objectives, this issue was categorized as an unexpected learning issue.

Procedure

Unexpected learning issues were examined regarding their relevance to the content of the course. A faculty member who was familiar with the course content (the same teacher used in Study 2, who was responsible for the design of the course) was asked to judge the relevance of these learning issues. He had to indicate whether an issue was irrelevant, fairly irrelevant, neutral, fairly relevant, or relevant. In addition, exploratory analyses were conducted to examine why students generated unexpected learning issues. For this purpose, the teacher was asked to sort these learning issues into meaningful categories.

Results

The 12 tutorial groups generated 520 learning issues for the 12 problems. The raters agreed that 32 learning issues (6.2%) definitely did not correspond to any of the faculty objectives. Of these 32 issues, eight were then judged fairly irrelevant, nine neutral, 14 fairly relevant, and one relevant. Consequently, 47% of the unexpected learning issues seemed to be at least fairly relevant to the course content.

The sorting procedure, which was conducted to identify key features of these unexpected issues, revealed four categories: (1) learning issues related to prior knowledge deficiencies; (2) learning issues focusing on patient management and medical intervention rather than on physiology or pathophysiology; (3) learning issues associated by faculty with additional curricular activities; and (4) learning issues arising from students' personal interests in, and experiences with, the subject matter.

The first category consists of learning issues associated with deficiencies in the students' prior knowledge. The faculty had not expected these learning issues to be generated because the issues were assumed to be a prerequisite to an initial understanding of the problem, studied by students in courses earlier in the curriculum. For example, a problem concerning the topic of child vaccination was developed to encourage students to study child vaccination schedules. Some tutorial groups, however, generated issues related to the immune system, although teachers assumed that this basic knowledge already had been mastered during a particular first-year course. Thus, the learning issues associated with the immune system were related to deficiencies in the students' prior knowledge.

The second category contains learning issues to be addressed in later years of the program. The first two academic years at the medical school of the University of Limburg emphasize normal physiologic and pathophysiologic processes. Patient management and medical intervention are addressed in the third and fourth academic years. The problem of the tall girl gave rise to learning issues such as "medical intervention of abnormal growth," "medication side effects influencing normal growth," and "abnormal growth rate due to an intracerebral tumor." These learning issues focus on diagnosis and therapy rather than physiologic processes. This implies that students sometimes generate learning issues that normally would be dealt with in later parts of the program.

Learning issues in the third cate-

gory are associated with additional curricular activities such as training in physical examination skills. Although these issues are presented to the students in the laboratory rather than through problems, some students did generate these issues. A possible explanation is that these students actually did not participate in the training. For example, one tutorial group generated the learning issues "examination of the reflexes of a newborn," whereas this topic could have been addressed during skills training. Another possible explanation is that skills training had raised students' interest in the subject matter, motivating them to spend more time on it.

The fourth category contains learning issues that can only partially be derived from the content of the course, because they were generated in response to salient details in the text of the problem rather than to the main focus of the problem. For example, a problem dealt with prenatal examination. One tutorial group generated the issue "purpose of physical training during pregnancy," referring to one sentence in the text of the problem in which the phrase "physical training during pregnancy" was mentioned. It was assumed that such learning issues arose from the students' personal interest or experiences. Furthermore, these issues are mostly connected with topics often discussed on television, newspapers, and other mass media.

In summary, it can be concluded that the students mainly generated unexpected learning issues in response to their own needs and interests. Although these issues were not planned by the faculty, 47% of them were rated either fairly relevant or relevant to the course content.

GENERAL DISCUSSION

The result of the present study — that an average of 64% of the faculty objectives for a given problem were actually generated as learning issues by the students — is nearly similar to the result reported by Tans et al.[12] They found an overlap of 68.3%. Both of these results differ from the percentage of overlap found by Coulson and Osborne.[10] They reported that each faculty objective was generated on average by five of 12 tutorial groups, for an overlap of 42%.

In the present study, the tutorial groups definitely failed to identify about 15% of the faculty objectives as being relevant to their learning. Analysis showed that these objectives were directly related to curricular activities other than specific problems, spanned more than one problem, or were psychological and social objectives rather than biological objectives.

The percentage of student-generated learning issues that were not expected (i.e., did not correspond to any of the preset faculty objectives) was 6.2%. Of these unexpected learning issues, 47% were judged by a teacher to be either fairly relevant or relevant to the objectives. Both of these percentages are low in comparison with the findings of Tans et al.,[12] who found that 21% of learning issues did not correspond to faculty objectives and that 77% of unexpected issues were either fairly relevant or relevant.

The results of the present study, including the nature of the unexpected learning issues (relating to students' deficiencies in knowledge of previous course work, anticipations of future course work, non–small-group curricular activities, and personal interests and experiences), suggest that students in a problem-based curriculum are able to determine what they need to know and what is relevant to learn. Further, students seem to modify their learning to satisfy their own needs and interests.

It should be noted that it is not always certain that students actually pursue the learning activities suggested by the learning issues they raise. Probably, students will not study certain issues related to a particular problem when they have already studied the relevant material. It is also possible that during self-study students spend time on issues that were not generated initially but, nevertheless, appear from their study of the literature to be related to the problem. Moreover, it is possible that students pay much attention to a specific issue during discussion in the tutorial group but eventually decide not to list this issue. Thus, student-generated learning issues define what students plan to study, but it is not certain that students will actually study the issues they generate. Learning issues may only to some extent reflect actual learning. However, the available evidence suggests that learning issues are reasonably valid indicators of students' learning, since Tans et al.[12] found a correlation coefficient of .42 between the frequency with which certain issues were listed and achievement on corresponding test items. Additional research is needed to further clarify the relationship between learning issues, activities, and outcomes.

The percentage of overlap between faculty objectives and learning issues provides opportunities to detect ineffective problems, i.e., problems that do not lead students to generate the appropriate learning issues. The degree of overlap can be used to evaluate the effectiveness of newly designed problems, for instance, during curricular innovations. In the present study, the overlaps between the faculty objectives and student-generated learning issues ranged from 27.7% to 100% for the 12 problems (Figure 1). The one problem that had an overlap of only 27.7% seems to have been an ineffective problem. Although examining overlap percentages seems to be useful in detecting ineffective problems, additional information about the nature of the specific shortcomings of the problems is required to improve them.

Problems to be designed for use in problem-based curricula should match the students' levels of previously acquired knowledge, should be concretely formulated, should refer to students' future professions, and should be briefly formulated. In addition, the problems should be open enough to stimulate discussion.[14] However, further research is needed to concentrate on the question of how to design effective problems. Despite the importance of problem effectiveness to guarantee coverage of a problem's content, little is known about the criteria by which problems could be designed or improved in order to guide students to preset faculty objectives.

References

1. Norman, G. R., and Schmidt, H. G. The Psychological Basis of Problem-based Learning: A Review of the Evidence. *Acad. Med.* **67**(1992):557–565.
2. Glaser, R. The Maturing of the Relationship Between the Science of Learning and Cognition and Educational Practice. *Learning and Instruction* **1**(1991):129–144.
3. Brown, J. S., Collins, A., and Duguid, P. Situated Cognition and the Culture of Learning. *Educ. Res.* **18**(1989):32–42.
4. Tulving, E., and Thomson, D. M. Encoding Specifity and Retrieval Processes in Episodic Memory. *Psychol. Rev.* **80**(1973):352–373.
5. Barrows, H. S., and Tamblyn, R. M. *Problem-based Learning. An Approach to Medical Education.* New York: Springer, 1980.
6. Schmidt, H. G., De Volder, M. L., De Grave, W. S., Moust, J. H. C., and Patel, V. L. Explanatory Models in the Process-ing of Science Text: The Role of Prior Knowledge Activation Through Small-Group Discussion. *J. Educ. Psychol.* **8**(1989):610–619.
7. Schmidt, H. G. Problem-based Learning: Rationale and Description. *Med. Educ.* **17**(1983):11–16.
8. Blumberg, P., Michael, J. A., and Zeitz, H. Roles of Student-generated Learning Issues in Problem-based Learning. *Teach. Learn. Med.* **2**(1990):149–154.
9. Blumberg, P., and Michael, J. A. Development of Self-Directed Learning Behaviors in a Partially Teacher-Directed Problem-Based Learning Curriculum. *Teach. Learn. Med.* **4**(1992):3–8.
10. Coulson, R. L., and Osborne, C. E. Insuring Curricular Content in a Student-Directed Problem-based Learning Program. In *Tutorial in Problem-Based Learning. A New Direction in Teaching the Health Professions,* H. G. Schmidt and M. L. De Volder, eds., pp. 225–229. Assen, The Netherlands: Van Gorcum, 1984.
11. Shahabudin, S. H. Content Coverage in Problem-based Learning. *Med. Educ.* **21**(1987):310–313.
12. Tans, R. W., Schmidt, H. G., Schade-Hoogeveen, B. E. J., and Gijselaers, W. H. Sturing van het onderwijsleerproces door middel van problemen: Een veldexperiment. [Directing the Learning Process by Means of Problems: A Field Experiment]. *Tijdschr. Onderwijsres.* **11**(1986):35–46.
13. Landis, J. R., and Koch, G. G. The Measurement of Observer Agreement for Categorical Data. *Biometrics* **33**(1977):159–174.
14. Majoor, G. D., Schmidt, H. G., Snellen-Balendong, H. A. M., Moust, J. H. C., and Stalenhoef-Halling, B. Construction of Problems for Problem-Based Learning. In *Innovation in Medical Education: An Evaluation of Its Present Status,* Z. M. Nooman, H. G. Schmidt, and E. S. Ezzat, pp. 114–122. New York: Springer, 1990.