

Active and Cooperative Learning in Signal Processing Courses

John R. Buck and Kathleen E. Wage

March 2005

IEEE Signal Processing Magazine,
vol. 22(2), pp. 76-81

© 2005 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Active and Cooperative Learning in Signal Processing Courses

How well do students learn the core concepts in signal processing courses? Do students learn better through some instructional formats than others? How can we assess student learning in these courses? This column describes our positive experiences using *active and cooperative learning* (ACL) methods to improve signal processing instruction. We provide examples, references, and assessment data that we hope will encourage other instructors to consider this approach. Many of our conclusions are based on impressions gathered through conversations with students during office hours as well as on responses from anonymous student opinion surveys. In addition to these subjective assessments, preliminary quantitative data measured with the Signals and Systems Concept Inventory (SSCI) support the benefits of ACL techniques in signal processing courses [1].

Our interest in ACL was sparked by an intriguing study on pedagogical methods in physics classes. A survey by Hake [2] of over 6,000 students in Newtonian mechanics classes found that students in traditional lecture courses learned only about a quarter of the concepts that they didn't know at the start of the course. In contrast, Hake found that students in *interactive engagement* (IE) format courses learned nearly half of the concepts that they didn't know at the start of the course. Hake defined IE methods as those promoting "conceptual understanding through interactive engagement of students in heads-on . . . and hands-on . . . activities which yield immediate feedback through discussion with peers and/or instructors." [2] The students' conceptual understanding

was assessed by administering the Force Concept Inventory (FCI) exam in a pretest/posttest protocol at the start and end of the course. The FCI, developed by Hestenes et al. [3], is a multiple choice exam emphasizing conceptual understanding of Newtonian mechanics over rote calculation.

Due to the work of Hestenes, Hake, and others like them, physics departments are increasingly adopting IE methods in their curricula [4], [5]. Over the past decade, the Foundation Coalition and other NSF-funded projects have championed engineering curriculum reform and assessment based on these pedagogical methods [6]. Still, relatively few advanced undergraduate courses employ IE classroom techniques.

Hake's survey gave us pause and caused us to speculate on the effectiveness of our own lecture-oriented courses. Like many instructors teaching signal processing courses like signals and systems and DSP, we believed that we emphasized the core concepts of the material in our lectures. Nevertheless, nearly all of our assessments were exams and homework that consisted predominantly of problems to solve rather than conceptual questions to answer. Moreover, we wondered if IE techniques, or ACL techniques as they are described in [7] and [8], would improve our students' understanding of signal processing concepts. These pedagogical methods incorporate a wide range of elements such as in-class problem solving, peer instruction, computer exercises, and interactive labs. All of these methods actively engage students in the immediate application of the key concepts of the course. This classroom format makes students responsible participants in their education rather than passive consumers

of lectures delivered by an instructor.

Since 1998, we have used some form of ACL in ten undergraduate and introductory graduate signal processing courses at the University of Massachusetts Dartmouth (UMD) and George Mason University (GMU). The following section presents concrete examples of how we use ACL in our courses. Subsequent sections discuss our insights on creating an effective ACL environment in the classroom, and our assessment data supporting the benefits of ACL in improving students' conceptual understanding.

ACL IN THE CLASSROOM

Our ACL courses incorporate interactive elements adapted from several sources. We found the books by Mazur [9] and Johnson et al. [7] helpful in developing ACL techniques for our signal processing classes. Mazur's book describes his implementation of concept-oriented peer instruction for his Harvard physics classes [9]. Johnson et al.'s book on cooperative learning provides several good ideas on formats and assessments for cooperative learning [7]. Other helpful ACL resources include [8], [10], and [11]. Our implementation of ACL has two major components: 1) reading quizzes to motivate students to prepare for class and 2) in-class problem-solving exercises to reinforce important concepts. Since accountability is crucial to the success of ACL methods [7]–[9], we grade the reading quizzes and in-class exercises; they are worth approximately 10% of the overall course grade. In the following paragraphs, we describe each of these ACL classroom elements in more detail.

Students spend the first five minutes of each class individually completing a four question true/false quiz on the assigned reading. These readiness assess-

ment tests, affectionately known as RATs, are graded to hold the students accountable for their own preparation for class. (See “Sample Readiness Assessment Test (RAT)” for an example of one of these quizzes on convolution and LTI systems). The purpose of these simple questions is to ascertain whether or not the students did the reading, rather than to assess conceptual understanding. After collecting the RATs, we verbally review them, asking students

to volunteer their responses along with their reasoning. This provides immediate feedback and also serves to outline the important concepts we will discuss that day. We use the RAT scores as one component of the course grade to stress the need to prepare for class by reading. To further emphasize the importance of the reading, we warn the students that the exams usually include a few true/false questions based on the RATs and then follow through on that warning. Other instructors have found the just-in-time-teaching strategy an effective alternative to RATs [11].

Collaborative problem-solving exercises are the central component of all of our ACL courses. We intersperse short (typically 10–15 minute) lecture segments on key concepts with 10–15 minute in-class group exercises. In-class exercises challenge students to apply what they have learned, both through reading before class and listening during the lecture segment, to solve a problem or answer a question. Students receive a grade for their answers to the in-class problems, providing an incentive for them to come to class well prepared to participate in group discussions. The students work on these exercises in groups, and at the end of the time allotted, the group is assessed based on one randomly chosen member’s solution. If time allows, the student may present his or her group’s solution orally. Groups not presenting their solution orally in class are assessed based on the written solution of a randomly chosen member. Both

SAMPLE READINESS ASSESSMENT TEST (RAT).

Readiness Assessment Test (RAT)

- 1) Name: _____
- 2) T/F Any discrete time signal can be written as a sum of delayed and scaled impulses.
- 3) T/F The convolution sum is a valid method of computing the output for any system, not just linear and time-invariant systems.
- 4) T/F Convolution *is not* an associative operation, i.e.,

$$x[n] * (h_1[n] * h_2[n]) \neq (x[n] * h_1[n]) * h_2[n].$$

- 5) T/F For a linear time-invariant system to be causal, its impulse response must be equal to zero for all time.

approaches fulfill our goal of making group members accountable to their peers. Grading these short exercises need not be onerous, and they provide valuable information to the instructor about the students’ understanding. Many feedback strategies can effectively promote group accountability and motivation so long as the feedback is prompt and performance based, not just a binary grade for effort. Other suggestions for assessment and accountability can be found in [7].

During the in-class exercises, the instructor circulates among the groups, observing the students’ progress and offering feedback or hints as necessary. The most productive hints are often posed as questions such as “If the system is causal, what can you say about the output here? Does your answer agree with that?” These questions often spark constructive discussions within the team. The students receive immediate and sometimes brutally honest feedback from their peers when discussing their proposed solutions within the group. Listening to the group discussions provides the instructor with immediate feedback on what the class understands and what they are confused about. In-class exercises can take many forms from simple pencil-and-paper problems to MATLAB programming.

SAMPLE SIGNAL PROCESSING ACL EXERCISES

As lecturers, we often solve several examples in class to illustrate key points. Just as professional athletes make difficult

plays look easy, professional signal processors can make complex examples look simple in lecture. While this is sometimes helpful, watching does not necessarily translate to understanding (or even to the ability to do the homework problems). Just as in athletics, watching a “pro” is no substitute for practice. When we simply lectured, our students sometimes turned in incomplete assignments with comments to the effect that while the class examples looked

easy, they got stuck or confused on the homework. We have found that the simplest way to engage students is to have the second in-class example for any concept be an ACL exercise. We solve the first example and then ask the students to work the second example in groups. For instance, during a class on the z -transform, we show students how to compute the z -transform of the signal $x_1[n] = (1/3)^n u[n-1]$ and then ask them to compute the transform of $x_2[n] = (1/5)^n u[n-2]$. This exercise allows students to practice some simple computational skills while also assessing whether or not they understand the basic idea of the region of convergence (ROC). It often reveals weaknesses in students’ math background (their facility with summations, in particular) that are easily addressed with an explanation from a peer or the instructor.

Convolution is also well suited to interactive instruction. One difficulty with working convolution problems is that they require time-consuming drawings of the signals. To speed up the process, we have the students draw the signals on transparency paper that can be flipped, shifted, and erased with ease. This low-tech solution dramatically helps the students who have difficulty visualizing the flipping and shifting operations. One of our students felt this approach was so helpful that he asked permission to bring blank transparencies to the exam. Solving convolution problems in groups helps students catch each other’s bookkeeping or arithmetic mistakes and increases the

likelihood that they obtain the correct answer. We believe many students receive a significant confidence boost by overcoming their initial mistakes in class with the support of their peers and the instructor. This makes them more likely to persevere when they encounter difficulties in the homework. Finally, these experiences reinforce the benefits of working in cooperative groups for the homework assignments.

In addition to straight-forward practice problems, in-class exercises can also challenge students to synthesize several concepts to answer a single question. “Example of In-Class Exercise” shows an exercise that requires students to make some general conclusions about the stability of FIR and IIR systems. The in-class problems that preceded this exercise provided several examples of both types of systems for students to draw upon. The definition of stability was presented in a previous class, so this exercise pushed students to link concepts from different classes. The discussions sparked by these types of questions force students to move beyond rote calculations and to think about what their results mean.

We have found that “reverse reasoning” exercises are very effective for improving students’ conceptual understanding. These problems reverse a typical class example by switching which quantities are the givens and which are the unknowns. If the first example provides the input spectrum $X(j\Omega)$ and the frequency response $H(j\Omega)$ and asks for a sketch of the output spectrum $Y(j\Omega)$, the subsequent reverse reasoning version of the problem might provide a different $X(j\Omega)$ and $Y(j\Omega)$ and ask which of several candidate system functions is consistent with that information. Reverse reasoning is also used in physics ACL classes, as discussed in [12]. We believe that students think more thoroughly about the concepts when solving these types of problems than when they are simply pattern matching new figures or expressions with a previous example. Reverse reasoning skills are crucial for engineering

EXAMPLE OF IN-CLASS EXERCISE.

In-class exercise: stability of FIR/IIR systems

- a) Choose the statement that best describes FIR systems.
- All FIR systems are stable.
 - Some, but not all, FIR systems are stable.
 - No FIR systems are stable.
- Give examples to support your conclusion.
- b) Choose the statement that best describes IIR systems.
- All IIR systems are stable.
 - Some, but not all, IIR systems are stable.
 - No IIR systems are stable.
- Give examples to support your conclusion.

design. Sherlock Holmes once noted

There are few people, however, who, if you told them a result, would be able to evolve from their own inner consciousness what the steps were which led up to that result. This power is what I mean when I talk of reasoning backwards . . . [13]

We want our engineering students to be among those few people who possess this power of reasoning backwards.

MAKING ACL WORK

There are several important issues associated with adopting ACL methods. First and foremost is the student reaction to this alternative approach. Many instructors find that students initially, and sometimes vociferously, resist the shift in format that requires them to be responsible for the material on a daily basis rather than simply show up for the course. It is important to acknowledge that this requires more work up front for the students, but also to explain that they will benefit from this based on our experience in previous courses and documentation in the literature. We make it clear that we will not be reverting to standard lectures. After roughly a quarter to a third of the term, almost all students accepted the new format; on the end-of-semester anonymous opinion surveys, we found that the majority of students ultimately preferred the ACL format course. Some students are quite enthusiastic; one student said that in-class problems are the “best teaching tool ever encountered.” Another commented that the exercises gave him “an opportunity

to spot areas of concern while the instructor was present so they could be addressed properly.” Another common student comment on ACL is that it makes class periods fly by quickly.

It is important to recognize that it takes time for students to adjust to the new method. Some students initially feel intimidated having to interact so much with their peers and the instructor. One said “The way you teach is much different than any other teacher

and at the beginning seems difficult and scary.” This student went on to note that as the semester progressed, he realized that his knowledge of signals and systems was increasing and he was grateful for that. At the end of the term, some students still feel strongly that the instructor should work all the examples on the board. While some may never be convinced of the benefits of active classes, one student said that he changed his opinion during the second ACL course he took, deciding that the new format was better. Compared to their GMU counterparts, most students at UMD accepted ACL class formats with much less resistance since most participated in a first-year engineering curriculum using an ACL format.

Interpersonal dynamics within teams is an important issue in ACL courses. Instructors may encounter forceful students who dominate a team or disengaged students who do not contribute to a team. These issues remain the most challenging in our ACL courses. Some colleagues cite these issues as their major reason for not adopting ACL or including group work in the students’ grades. We believe that some healthy friction and challenge in team interactions is an important element of the students’ education for the “real-world” conditions of their future engineering careers. Most students will participate in project teams where their personal performance reviews include an element of the team’s performance. The depiction of the engineering workplace in Scott Adam’s *Dilbert* is drawn from life, all too often equal parts documentary and comedy. Dysfunctional teams are real issues

in engineering projects. As responsible instructors, we owe our students some supervised classroom experiences navigating the challenges of team dynamics before their professional success depends upon it. A number of helpful suggestions for addressing these concerns can be found in [7]. Ultimately, our courses are structured so that even extreme issues in group assignments do not cause a student to lose or gain more than half a letter grade. We believe this risk is justified by the benefits of the ACL experience for the majority of the students.

Another common concern of instructors is that the time given to in-class exercises will not allow them to cover all the material on the syllabus. However, as we often remind ourselves, “Just because you said it, doesn’t mean they got it.” Our preliminary assessments mirror those of Hake in [2] (see the section below). Even when everything is “covered” in a traditional lecture course, assessment indicates that the students are only learning about a quarter of the concepts. Rather than lecturing on all the material, we believe it is better to spend class time on interactive exercises that reinforce the core concepts. Studies have shown that students who have a conceptual framework for organizing information can apply what they have learned to new situations and can learn related information more quickly [14]. We help students to develop this conceptual framework in class and then expect them to learn additional material outside of class through homework assignments. To reinforce the idea that students are responsible for all the reading, not just what is covered in class, we assign homework and exam questions on material included in the reading that is not discussed in class. We make it clear to students that this will be the case (and after the first exam contains a problem based on one of these homework problems, most take us seriously). The course description and outcomes remain unchanged even if we don’t explicitly talk about every aspect of every topic each time we teach the course; rather, we prioritize those topics that the students find confusing or that we believe to be difficult to understand.

The success of the ACL format depends on students reading the assigned textbook material before class. In our initial experiences with ACL, we were very concerned that our students would not complete the reading assignments. We were pleasantly surprised that the daily reading quizzes provided sufficient incentive for most students to prepare for class. Interestingly, many students are also in favor of the RATs. The RATs were cited by 40% of GMU students on anonymous end-of-term surveys when asked an open-ended question concerning what should not change about the class. One student commented that the RATs are “tough love and should continue.” Statistical analysis of RAT scores at both schools indicates that most students perform significantly ($p < 0.05$) better than chance, suggesting that they are reading in preparation for class.

Many instructors considering the adoption of ACL methods find it intimidating to relinquish the control that a lecture format gives them. Adopting one or two ACL elements in a course may be less threatening for both the students and the instructor, and still provide some of the benefits of ACL instruction. The introduction of ACL methods into courses at GMU began with assigning a collaborative warm-up question at the start of each traditional lecture class while students arrived. These warm-ups were short exercises intended to refresh students’ memories about concepts from the previous class and to prepare them for the material to be covered that day. Often, the exercises required a short calculation that was later used in an example in the class. We encouraged students to discuss problem-solving strategies and compare answers with their peers. On the end-of-term surveys, student response to the warm-up questions was overwhelmingly positive. As we grew more comfortable with the method, we introduced more ACL elements into our courses until we converged on the class format described above. We found that it takes less time to prepare ACL classes than lectures, but that a good ACL course requires more attention and energy from us during class; we must think on our feet and engage the students in the material.

ASSESSING STUDENT LEARNING

Both our pedagogical instincts and the student opinion surveys told us that the students were learning more when we adopted ACL techniques in our courses. Our experiences, like those reported in [15], suggest that these surveys must be interpreted with a grain of salt, as positive student reviews do not guarantee that students are learning. We wanted to confirm these subjective assessments with an objective quantitative measure of student learning like the FCI. Since no analog to the FCI existed for signals and systems, we developed the SSCI. The SSCI is a 25-question, multiple-choice exam designed to be given in one hour, and is available in both discrete-time (DT) and continuous-time (CT) versions. The exam tests the core concepts of linearity, time-invariance, transform representations, convolution, filtering, sampling, and background mathematics in a manner emphasizing conceptual understanding over mechanistic problem solving. If the students understand the concepts tested, they can choose the correct answer without performing any written calculations. Conversely, if the students do not understand the concepts tested, the questions contain little or no information that they can plug into rote calculations or memorized formulae. Also, each question’s incorrect choices, or distractors, embody common student misconceptions about signals and systems.

The current versions of the SSCI are the culmination of three years of development and calibration described in [1]. To date, 28 instructors at 12 institutions have administered the SSCI in some form to over 1,000 students in signal processing classes. The SSCI has been used both for instructors’ personal assessment of student progress and also as one component of departmental ABET assessment strategies. Qualified instructors can download the most recent versions from the SSCI Web site [16] after e-mailing the authors to obtain the passwords used to protect the exams. The paper cited above describes several interesting results from our pedagogical study using the SSCI. For this column, we will focus on using the SSCI to assess instructional techniques.

Administering an assessment exam like the SSCI at the beginning and end of a course and computing the difference between the pretest and posttest scores is a standard technique used to quantify how much students have learned [2], [17]. If the exam is only administered as a posttest, it assesses achievement, rather than improvement, since there is no control for the students' initial knowledge. As in Hake's survey [2], we use normalized gain $\langle g \rangle$, defined as

$$\langle g \rangle = \frac{\text{posttest} - \text{pretest}}{100 - \text{pretest}},$$

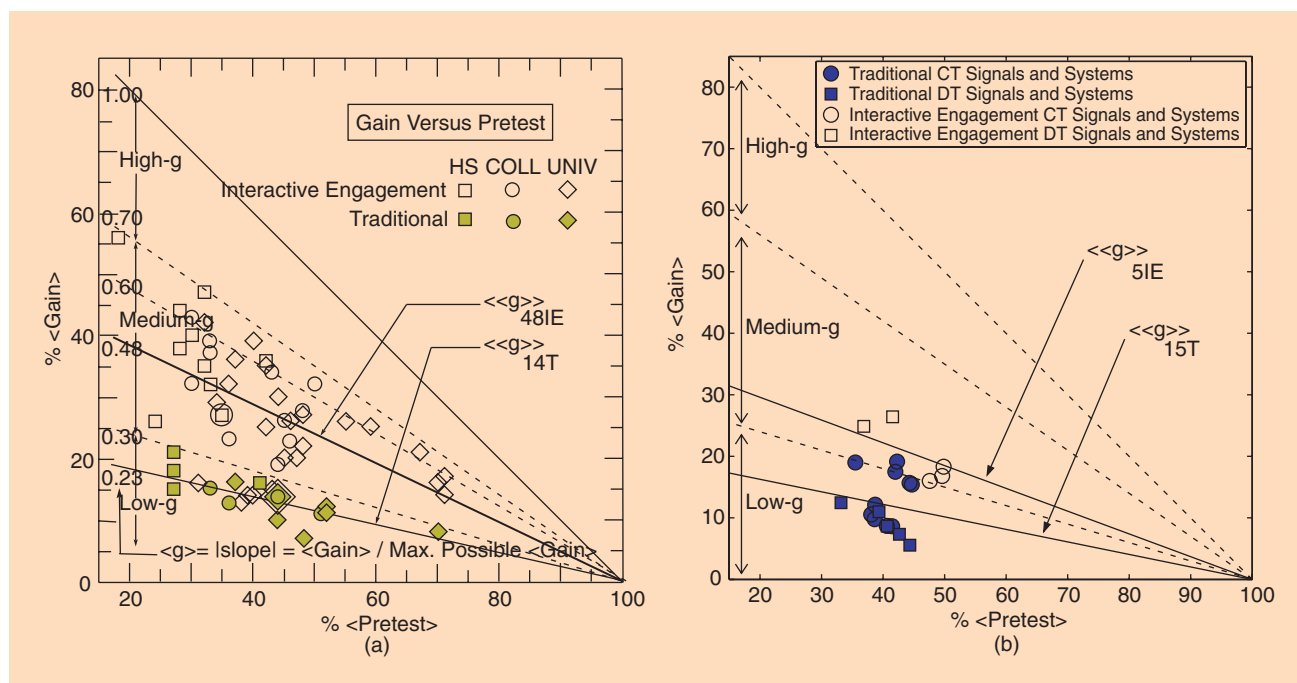
as a metric for the improvement in students' conceptual understanding. The pre- and posttest values are the averages computed for the course using only the cohort of students who took both the pretest and posttest. Thus, normalized gain is the fraction of the available improvement in score that was achieved

in the course. An equivalent interpretation is that the students learned $100\langle g \rangle\%$ of the material that they didn't know at the start of the course.

Pedagogical research in physics has found that $\langle g \rangle$ is robust to variations in instructor experience, student background, class size, and university ranking [2], [3]. Hake's major conclusion was that 14 traditional lecture format classes achieved normalized gain $\langle g \rangle = 0.23 \pm 0.04$, while 48 IE (or ACL) courses achieved $\langle g \rangle = 0.48 \pm 0.14$, nearly two standard deviations better than lecture courses. Subsequent papers have reported similar performance for IE methods in physics courses [4]. In our study using the SSCI, we found results strikingly similar to those reported by Hake. We computed $\langle g \rangle$ for 20 signals and systems courses. The 15 lecture format courses had normalized gain $\langle g \rangle = 0.20 \pm 0.07$, while the five ACL courses for which we have data achieved $\langle g \rangle = 0.37 \pm 0.06$. The gain for these

ACL courses is more than two standard deviations above the lecture courses.

Figure 1 contains two plots illustrating the similarity between our results and those in Hake's survey [2]. Figure 1(a) is reproduced from [2] by permission of Richard Hake and the American Physical Society, while Figure 1(b) contains data measured using the SSCI. Both figures plot the class average pretest score on the abscissa and the class average raw gain (posttest minus pretest) on the ordinate. Using these axes, lines of constant $\langle g \rangle$ appear on the graph as radiating out of the bottom right corner with a constant slope. The average values of $\langle g \rangle$ for each class format are represented by solid lines labeled with $\langle\langle g \rangle\rangle$. In Figure 1(a), the 14 yellow shaded points represent the lecture format classes, while the 48 open points represent IE format classes. The different symbol shapes indicate whether the data point represents a high school (square), college (circle), or



[FIG1] Comparison of the Force Concept Inventory (FCI) and Signals and Systems Concept Inventory (SSCI) results contrasting students' gain in conceptual understanding for IE/ACL format courses versus lecture format courses. Each point represents a single course, with the abscissa being the average pretest score and the ordinate being the raw gain. Gain is defined as the posttest average minus the pretest average (*i.e.*, $\% \langle \text{Gain} \rangle = \text{posttest} - \text{pretest}$). Both exams find that students in IE/ACL courses (open points) achieve a normalized gain that is roughly two or more standard deviations above the normalized gain in traditional lecture courses (shaded points). (a) FCI results: Comparison of the effects of traditional (T) and interactive engagement (IE) pedagogical methods on students' conceptual understanding of Newtonian mechanics as assessed by the FCI [3]. The plot shows data for 6,542 students in 14 traditional and 48 interactive courses. (Reprinted from [2] with permission). (b) SSCI results: Equivalent gain data for signals and systems courses assessed using the SSCI [1] in a pretest/posttest protocol. The plot shows data for 600 students in 15 traditional and 5 interactive engagement courses.

university (diamond) class. Note that in Hake's data all courses in the medium gain region ($0.3 \leq \langle g \rangle \leq 0.7$) are IE courses. Figure 1(b) plots the data from the SSCI in a similar format. Here, the blue shaded points represent traditional courses, while the open points represent IE courses. Circles are used to represent CT signals and systems courses and squares are used to represent DT ones. In this limited data set, the IE courses achieve larger values of $\langle g \rangle$ than the traditional courses. All five ACL classes fall in the medium gain region ($0.3 \leq \langle g \rangle \leq 0.7$), while only two of the 15 traditional lecture courses achieved a $\langle g \rangle$ in this region. The similar values of $\langle\langle g \rangle\rangle_T$ obtained for the two exams using a comparable number of traditional courses encourages speculation that normalized gain in lecture classes is not only robust to instructor experience, class size, and student background, but perhaps also to the material taught. The smaller data set on ACL methods for the SSCI precludes strong conclusions, but the congruence of this limited data with that reported for physics classes ([2], [4], [5]) supports our subjective assessments that ACL methods improved our signals and systems courses.

CONCLUDING REMARKS

Our data show that our students learn core concepts in signal processing better when the class requires active participation than when a traditional lecture format is used. We strongly encourage other instructors to try ACL methods in their signal processing courses. We've provided some samples from our courses to spark ideas. Additional ACL information and examples can be found at [6] and [16]. Also, it is not hard to adapt textbook and exam problems into class exercises. Nothing motivates student interest like a casual comment that the current exercise was taken from last year's final exam.

We hope that instructors who are using ACL will share both their experiences and materials. Finally, we encourage instructors of both ACL and lecture courses to use the SSCI in a pretest/posttest protocol in their courses to

measure how much their students are learning. Some of the leading pedagogical reformers in physics were skeptical of Hestenes et al.'s initial results until they administered the FCI in their own classes [15], [9]. If possible, send us the average pretest and posttest scores for your students so that we can further populate Figure 1 and determine whether our preliminary findings on the advantages of ACL hold up with a more extensive data set. To ensure good quality data for the study, we ask instructors submitting their class scores to follow the instructions and complete the survey on the SSCI Web site [16].

ACKNOWLEDGMENTS

The authors acknowledge the support of NSF grant EEC-9802942 to the Foundation Coalition for development of the SSCI. Laurie Fathe (GMU) and Nick Pendergrass (UMD) provided valuable support and mentoring in our adoption of ACL methods. Richard Hake has encouraged our efforts to assess the benefits of ACL methods for signal processing courses. Workshops by P. K. Imbrie, César Malavé, and Jim Morgan inspired our first ACL efforts. Laurie Fathe, Nick Laneman, Nick Pendergrass, and two anonymous reviewers made helpful suggestions on earlier versions of this article. We are indebted to instructors, too numerous list here, who contributed their SSCI data to Figure 1.

AUTHORS

John R. Buck (johnbuck@ieee.org) is an associate professor jointly appointed to the Department of Electrical and Computer Engineering and the School for Marine Science and Technology at the University of Massachusetts Dartmouth. He received his Ph.D. from the MIT/WHOI Joint Program in 1996. He was awarded the MIT Goodwin Medal in 1994 for conspicuously effective teaching by a graduate student. He is the co-author of several textbooks on signal processing. His research focuses on signal processing and underwater acoustics, and has been supported by an NSF CAREER Award and an ONR Young Investigator Award.

Kathleen E. Wage (k.e.wage@ieee.org) is an assistant professor in the Department of Electrical and Computer Engineering at George Mason University (GMU). She received her Ph.D. in electrical engineering from the MIT/WHOI Joint Program in 2000. As a graduate student, she won the Harold L. Hazen Teaching Award from the MIT EECS Department. In 2004, she received an Outstanding Teaching Award from GMU's School of Information Technology and Engineering. Her research focuses on array processing and underwater acoustics. She received an ONR Ocean Acoustics Entry-Level Faculty Award in 2002.

REFERENCES

- [1] K.E. Wage, J.R. Buck, T.B. Welch, and C.H.G. Wright, "The signals and systems concept inventory," *IEEE Trans. Edu.* (forthcoming).
- [2] R.R. Hake, "Interactive-engagement vs. traditional methods: A six-thousand-student survey of mechanics test data for introductory physics courses," *Am. J. Phys.*, vol. 66, pp. 64–74, 1998.
- [3] D. Hestenes, M. Wells, and G. Swackhamer, "Force concept inventory," *Phys. Teacher*, vol. 30, pp. 141–158, Mar. 1992.
- [4] C.H. Crouch and E. Mazur, "Peer instruction: Ten years of experience and results," *Am. J. Phys.*, vol. 69, pp. 970–977, 2001.
- [5] J.W. Belcher, "Improving student understanding with TEAL." [Online]. Available: <http://web.mit.edu/jbelcher/www/fnlEditedLinks.pdf>
- [6] Foundation Coalition [Online]. Available: <http://www.foundationcoalition.org>
- [7] D. Johnson, R. Johnson, and K. Smith, *Active Learning: Cooperation in the College Classroom*. Edina, MN: Interaction Book, 1998.
- [8] R.M. Felder and R. Brent, "Cooperative learning in technical courses: Procedures, pitfalls, and payoffs," ERIC Report 377038, 2004 [Online]. Available: <http://www.ncsu.edu/felder-public/Papers/Coopreport.html>
- [9] E. Mazur, *Peer Instruction: A User's Manual*. Englewood Cliffs, NJ: Prentice Hall, 1997.
- [10] "The jigsaw classroom: A cooperative learning technique" [Online]. Available: <http://www.jigsaw.org>
- [11] G. Novak, E. Patterson, A. Gavrín, and W. Christian, *Just-in-Time Teaching: Blending Active Learning and Web Technology*. Englewood Cliffs, NJ: Prentice-Hall, 1999.
- [12] A.V. Heuvelen, "Millikan lecture 1999: The workplace, student minds, and physics learning systems," *Am. J. Phys.*, vol. 69, no. 11, pp. 1139–1146, 2001.
- [13] A.C. Doyle, *A Study in Scarlet*, 1887 [Online]. Available: <http://www.gutenberg.net/etext/244>
- [14] J.D. Bransford, A.L. Brown, and R.R. Cocking, Eds., *How People Learn*. Washington, DC: National Academy, 2000.
- [15] E. Mazur, "Qualitative versus quantitative thinking: Are we teaching the right thing?," *Opt. Photon. News*, vol. 3, pp. 38, 1992.
- [16] Signals and Systems Concept Inventory Web Site [Online]. Available: <http://www.signals-and-systems.org>
- [17] R.R. Hake, "Assessment of physics teaching methods," in *Proc. UNESCO-ASPEN Workshop on Active Learning in Physics*, Univ. of Peradeniya, Sri Lanka, 2–4 Dec. 2002.