ICED 2020 proceedings:

Instructional design based on the problem-solving decisions of scientists and engineers

Carl Wieman[1]

Stanford University, Department of Physics
382 Via Pueblo Mall
Stanford, CA 94305

Argenta Price[2], **Candice Kim**[3]

Stanford University, Graduate School of Education
485 Laseun Mall
Stanford, CA 94305

Abstract

We have explored the detailed process by which skilled scientists, engineers, and doctors ("experts") solve authentic problems—the problems they are called upon to do in their work. Such problems are far more complex than typical problems used in university courses and exams. We find there is a set of 29 decisions that experts make in the solving process, a set which is nearly universal across the 31 experts and 10 disciplines we examined. These decisions provide a very useful framework for measuring problem-solving expertise. Here we focus on how they can also be used to design instruction that will improve students' problem-solving skills. This instruction gives them practice at making many such decisions in realistic contexts, a necessary step to learning to be good problem solvers, and hence, good scientists, engineers, and physicians. Practice they do not get in traditional instruction.

1 Introduction: The nature of expert problem-solving in science and engineering

We have been studying how skilled practitioners ("experts") in science and engineering solve authentic problems in their work. These are problems with far more complexity than have been used in previous research on problem-solving and expertise, or that students typically see in their courses. Such problem-solving skill is arguably the most important goal of education in science and engineering, but little is known about the details of the problem-solving process as used by experts, or how best to teach it. We have been able to characterize the problem-solving process in detail for the first time (Price, et al., 2020) in terms of the set of specific decisions experts make in the decision process. This provides a new and valuable framework for both teaching and measuring expertise in authentic problem-solving. Here, after introducing the basic concepts, we will discuss how to design instructional activities that will more effectively teach the authentic problem-solving skills that will be needed by students when they enter the technical workforce. Such activities require the student to explicitly practice making these problem-solving decisions used by experts, with guiding feedback.

In previous work (Jones, Madison & Wieman, 2015), one of us (CW) discussed the general design and classroom implementation of effective "active learning", built around the idea of having students carry out "deliberate practice" (Ericsson, 2018) learning in the classroom. This

---

[1] cwieman@stanford.edu
[2] argenta@ stanford.edu
[3] candicekim@stanford.edu

involves the students practicing and getting feedback on the specific cognitive actions used by experts in the discipline. This approach has been shown to be highly effective in teaching physics from the most introductory courses (Deslauriers, Schelew & Wieman, 2011), to advanced undergraduate (Jones, Madison & Wieman, 2015), to the most advanced graduate (Lepage, in press) courses. It has also been shown to be effective in many other disciplines and by many other authors. (In many cases they will use different labels, but the concepts are the same.) While the Jones, Madison, and Wieman (2015) reference discusses the details of implementation and the general principles for the design of the instructional activities, it leaves it to the instructor to figure out what the “expert” reasoning processes are that the student needs to practice and learn, and hence what the specific questions or tasks are that students should complete.

This paper provides further guidance to instructional design by identifying a detailed and specific framework for those expert reasoning processes. At the time of the Jones 2015 work, we never thought that such a framework was possible, that the details of the problem-solving process would have some very general characteristics in common, but most of it would be somewhat dependent on the individual and quite dependent on the discipline and topic. In recent research we carried out a detailed analysis of the problem-solving process of 31 experts spanning 10 disciplines of science, engineering, and medicine used in solving authentic problems in their discipline. We analyzed the process in terms of the specific decisions they made. We discovered, to our surprise, that there is very little variation in terms of the decisions. There is a consistent set of 29 decisions that these, and likely nearly all, experts in science and education make in solving problems, and a consistent way they organize their knowledge for making those decisions. Here we show how these decisions can provide much more detailed guidance to make the design of educational activities easier and more effective. This design provides students with practice and feedback in explicitly making these expert problem-solving decisions in realistic contexts.

Here the label “expert” refers to a skilled practitioner, such as the typical professor in science and engineering at a good research-intensive university or the typical experienced engineer or scientist at a well-established company or government research laboratory. In other words, high level but not extraordinary performers in the technical workforce. In our work, we have also included experts in medical diagnoses. Authentic problems are ones these experts solve as the core of their work, and involve research, design, and diagnoses. Such problems involve a high level of complexity, substantial investments in time, and have no clear route to a solution, and at the beginning it may not be clear that there even is a solution. These are characteristics that are missing from nearly all problems that university students will encounter in their coursework. The set of 29 decisions that our sample of experts used consistently across problems in science, engineering, and medical diagnosis included answering such questions as: “What are the most important factors in the situation?”; “What simplifications or approximations are appropriate?”; “What information is needed to answer this question?”; “How credible in this information?”; “What conclusions can be justified by this information?”; and “Does this solution satisfy all specific and general criteria for a good solution? What are potential failure modes, either literal or logical?” In traditional S & E course homework and exam problems, seldom do students have to answer questions like this, and for that reason they are often not prepared to perform as well as they might in the S & E workforce.

We find that the experts nearly always call on specialized disciplinary knowledge to make these decisions, and that knowledge is organized in a particular way. It is in the form of “predictive frameworks”, a mental model that incorporates all the key information and concepts and the relationships between them in terms of underlying mechanisms, while excluding unimportant details. Such a framework allows the expert to run mental simulations as to what behavior would be expected under different conditions, and how that behavior would change when various parameters in the model are changed. It also organizes the information into “chunks”, thereby reducing demands on the working memory, as has been previously observed in expert-novice comparisons (Ericsson 2006).

As a side note, our work argues strongly that there is no such thing as “critical thinking skills” that transcend all disciplines, as some have claimed (Arum 2009). While it is essential that a person recognize and be able to make all of these 29 decisions, if they are to be an expert problem-solver, actually making any of these decisions in an expert fashion requires disciplinary knowledge. Recognizing what knowledge is relevant and applying that knowledge appropriately in the context are important learned expert skills. Even if a person knows they need to make a decision such as “What factors are important?”, they would never be able to make such a decision correctly without having substantial disciplinary knowledge and being able to decide which information would apply in the specific context given. Tests that claim to measure general “critical thinking skills” independently of any disciplinary knowledge, such as the Collegiate Learning Assessment, are fundamentally flawed in their failure to recognize that the correct application of knowledge is an essential element of solving any authentic problem and hence of any meaningful “critical thinking”. Thus, any meaningful test of critical thinking will be discipline specific and the results will depend strongly on whether or not the testee knows the relevant content.

2 Assessing expert problem-solving

This framework of decisions and predictive frameworks has given us a template for designing assessments of student problem-solving skills that can be readily mapped onto different disciplines. We explain this in some detail, as it is closely related to the template used for instruction. The primary goal of these assessments is not to rank students individually, as in the typical course. Instead it is to compare every student with an absolute standard—how similar are they to how a skilled practitioner solves a realistic problem? While this can be used for feedback to individual students as to how they can improve, more often it is used to look at the aggregated results from an entire class or program to determine how effectively these students are being educated to perform like skilled practitioners, and where the strengths and weaknesses are in their education. The assessment gives a realistic problem context and then has students make the decisions we have identified that are used by experts. We then examine how well their decisions and the reasoning behind each decision matches with an expert’s.

The assessment template involves first deciding on the realistic challenge: in medicine, a sick patient they need to diagnose; in engineering, a design problem; or in science, a phenomenon to explain or control. The assessment will then have the student answer a series of questions that correspond to particular problem-solving decisions from the expert decision list. The choice of the particular challenge context and the details of the questions determine what disciplinary knowledge is required and the level of complexity of the question, i.e. how many different factors need to be taken into account. Those are adjusted to be suitable for targeting particular courses or overall program learning outcomes. In some cases, to make the assessment more constrained and hence more practical due to constraints on time or scoring complexity, the problem will involve “trouble-shooting”, where the student is given a trial solution which contains multiple deficiencies by design. They are put in the position of a manager evaluating and improving the design of an intern. Consistent with the decision set and authentic problem-solving, at different points students will be asked what information they need to solve the problem, and at other times during the assessment they will be provided with new information to see how they can evaluate and use that information.

Such assessments are extremely revealing, showing dramatic differences between experts and students, including serious weaknesses in many areas that students have supposedly been taught. The results demonstrate the profound difference between learning particular content, items of information or calculational procedures that are probed by the typical course exam and recognizing how and when to use that knowledge to make decisions in solving problems. Across an increasing number of different disciplines (medicine, mechanical engineering, chemical engineering, earth sciences, physics) we see large differences between expert and student responses, both in the decisions they make and the rationale they give for those decisions (which are also part of the assessment). Although typical, this is not universal. Students' scores on the assessment questions are strongly correlated with what instruction they have received, and higher on the assessment questions where their instructional activities involved making decisions of that specific type.

3 Teaching expert problem-solving: Practicing making decisions

In this case the teaching is nearly identical to the assessment. In both cases students work through a realistic problem, making decisions and justifying them in the process. Their responses are compared with the generic responses that a skilled practitioner would give. The only real difference is the conditions in which the student is working. Rather than being isolated and without any feedback, as in the assessment, in the teaching they are regularly interacting with other students and the instructor to get feedback, new information, and assistance to support their learning.

Teaching a student to be an expert problem solver in a discipline requires them to practice making these problem-solving decisions, with guiding feedback to improve their decision-making capabilities. That implies they must also learn the relevant disciplinary knowledge in the form of predictive frameworks that are necessary in making such decisions. Thus, the design of effective instruction has a large amount of overlap with what we have described for the authentic assessment. For students to learn to solve authentic problems they need to work through a problem in an authentic context, practicing making the various problem-solving decisions experts make. As discussed in Holmes, Keep, and Wieman (2020), the most effective scaffolding is provided by telling them what decisions they need to make, but leaving it to them to make the decision, followed by reflection and feedback on their decision.

The design and implementation of the most effective instructional activities needs to include all elements of teaching expertise as discussed by Wieman (2019). This involves several elements not discussed here, including the need to have the activities at the correct level to connect with and build on the learners' past experiences; use a context and format to motivate learners to work hard to master the desired learning; and respect the limitations on the brain's working memory and so avoid all unnecessary cognitive load.

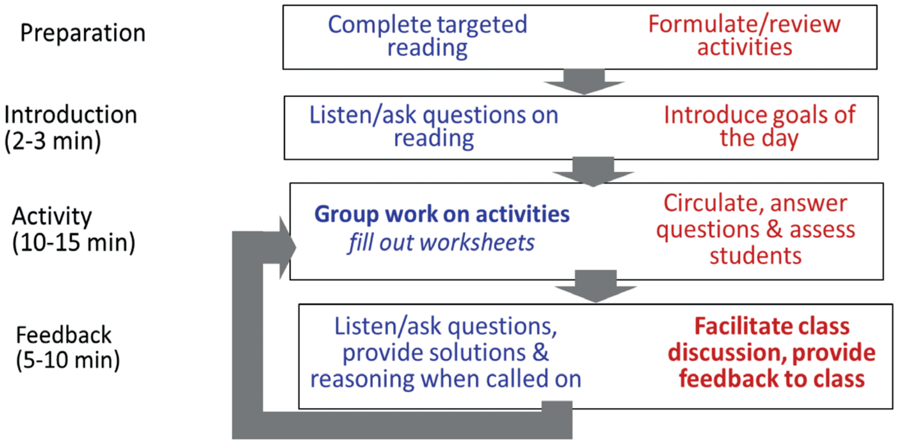

Figure 1. Structure of instructional activity, with very approximate timing to be used in class. The last two boxes will be repeated with new questions until the activity is completed. From Jones, Wieman, and Madison (2015)

In terms of implementation, the practice of and feedback on the instructional activities follows standard active learning teaching methods described in Wieman (2019) and by Jones, Madison, and Wieman (2015) (see Fig. 1). Students have individual deliverables, usually worksheets they must individually complete and turn in that involve answering a number of challenging questions that reflect decisions they must make and justify. The pre-class preparation will also often be the same as described there, but it may vary with the discipline, task, and instructional level. Students do some individual work first in class thinking about the questions, and then work as part of a small group to figure out answers that are too difficult for them to do by themselves. The group also provides additional information and feedback on the reasoning of individuals. The instructor monitors their efforts by listening in on these conversations, asking probing questions to individual groups, as well as by answering questions from groups and providing new ideas and information as needed. Most importantly, the instructor will provide timely specific feedback on the work. This will help guide their thinking by explaining where and how their reasoning is incorrect and how to change it (Schwartz, 2016a). This feedback is provided at the level of the individual, the group, and the class, as appropriate (Schwartz, 2016b).

This feedback also includes “just in time telling” (Schwartz, 2016b), where students are provided with new information/content, often through short lectures or sometimes readings, only after they have struggled with the problem, and then recognize they needed additional information and ideas to make progress. In this way, they are first “prepared for future learning” (Schwartz & Martin, 2004), and so are much more likely to structure the material they learn in terms of a predictive framework such as an expert would use, organized around how and when that material is useful for solving problems.

3.1 Components of the design of the instructional problem

Here we describe a general design for creation of the questions for students to answer. The instructional design is based around giving students explicit practice with feedback on making the set of decisions an expert would make in solving the problem. Although it seems almost unnecessary to say it, our research (currently in the process of being written up and published) supports common sense. Across multiple disciplines, we see that students get more expert-like in the decisions they practice making in a realistic disciplinary context in courses or other relevant experiences, and they remain very novice-like in how they make those decisions they have not practiced, even though they have covered all the content needed.

First, the students are given an authentic context and the overall goal of the task. Examples are: diagnose patient illness; a mechanical design to accomplish a particular purpose and relying on a particular technology; design of a chemical production process plant; figure out how strong the components of a structure like a bridge must be, etc. You may find it advantageous in some circumstances to give students a potential solution or design to troubleshoot, rather than to figure out a solution on their own with no guidance. This can constrain the solution space in pedagogically useful ways. Many of the expert decisions involve evaluating and improving potential solutions, and so such “trouble-shooting” problems can allow authentic practice of many expert decisions. They allow instructional activities to focus on these decisions in cases where there is insufficient class time for students to carry out a full solution, or they carry out such a range of different possible solutions that it would be difficult to keep the class coherent.

Below we list the types of questions students would be asked to complete. Where it seems useful, we list specific examples of what these might be in particular disciplines. In keeping with the universality of the set of expert decisions, many need no examples because they are so obviously applied everywhere, such as thinking of related problems one has seen previously and using information to decide between alternative solutions. Below that we list the set of decisions from the list of 29 that are being probed.

1. What features are the most important? (includes what information and concepts are relevant and irrelevant, how to best visualize the context). *This could be asked in the contexts of diagnosis of a patient in medicine or a malfunctioning science apparatus or experiment, or a possible engineering design. In these cases, as in the real world, students would be give a rich set of information, some would be relevant and some not. Students* *would be asked to identify what of the information available was most important for solving the problem (or critiquing the design or potential diagnosis), and what further information did they need, and why.*

2. What previous problems have you seen that are similar and relevant? What insights for solving this problem can you get from them?

3. What simplifications or approximations can you make, and why? *E.g.: In a science context this would normally involve a variety of factors that are neglected in the process of choosing simplified but accurate, and hence most useful, models of the problem context. In engineering design, it would involve recognizing what elements of the design are unimportant.*

4. What are tentative possible solutions/diagnoses/designs? *Students will likely require some guidance in answering this question as an expert would, namely roughing out possible approaches and making basic evaluations of feasibility, without attempting to carry out a detailed solution.*

5. How would you go about deciding on which solution is best? *E.g.: Experts carry out detailed plans at this point, laying out the solution process they plan to follow, including noting the greatest uncertainties and setting priorities, most students spend very little time on this step. As appropriate to the circumstances, students should be called upon to lay out a detailed plan. If they need scaffolding to help learn how to do this, provide worked examples, and possible plans for them to compare and critique to help them learn the important elements that should be included in a good plan.*

6. What additional information do you need to decide on a solution, and why? How will you use this information? *This is really a subpart of the planning, but you may or may not want to lay it out as a separate step to give it emphasis.*

Provide additional information to students. (This should include relevant and irrelevant information to give them practice at distinguishing between the two, as the real world is always filled with both kinds.)

7. Interpret and apply this information to refine your solution.

8. Narrow down your set of possible solutions.

9. Reflect on the problem-solving process you have followed and decide whether it is productive and will likely lead to a good solution, or whether you need to step back and think how the problem-solving process needs to be modified to be more effective.

Possibly provide additional information to students, targeting more specific features of likely solutions.

10. Which is the best solution, and why?

11. How can you test this solution based on disciplinary-specific criteria, and how well does it pass those tests?

This list is for illustrative purposes only, and to keep it reasonably short it does not include all of the decisions that a student needs to learn to make in solving problems in a particular discipline. Notably, it did not include any decisions related to the collection, evaluation, and application of data. For many types of problems those are very important, but they can be easily inserted into the general template that we have given here, as desired. Depending on the instructional context and goals, particularly the time available and the background experience of the students, you may want to ask only a subset of the questions we have listed, or add or substitute other questions from the full set of 29 listed in Price, et al (2020). Almost all of these are easy to include just by making those decisions items the student has to complete in the context given.

The problem described required decisions to be made on:

- the key features
- what prior solutions can be applied
- what simplifications and approximations are relevant and justifiable
- potential solutions
- plan for determining the solution
- plan for gathering information to test between possible solutions
- how valid/reliable information is
- how the information compares with expected results
- what conclusions are appropriate based on the data
- how well the problem-solving approach is working
- if previous decisions about simplifications and predictive frameworks need to be modified
- the best solution
- how well the chosen solution holds

4 Conclusions

What we describe above does not include teaching all the elements a person must master to be a skilled practitioner in these fields. First, because of the instructional constraints, the thoughtful instructor will need to make decisions and tradeoffs depending on the circumstances. This will make it impossible to cover the full set of decisions multiple times with different types of authentic problem, as is likely necessary to achieve true mastery. We would argue that by the time students are at an intermediate and advanced undergraduate level, students need to be practicing nearly all these skills in the context of their fields of specialization. The only exceptions are the few involving choice of problem, dissemination, and larger relevance. Arguably, graduate students need to be practicing and learning the entire set. Second, there are a small set of skills needed in the workforce that are not reflected in our list of problem-solving decisions, most notably social skills, such as working effectively as part of a team and elements of communication not covered in the decision list. These should not be neglected. Instruction designed around practice with feedback on making the set of expert problem-solving decisions may not teach everything you want students to learn, but it will cover far more than they are likely learning now. In work in progress we have implemented this approach in three different courses covering a very wide range of levels and disciplines, and preliminary results look good. Such instruction will provide unique and important preparation for becoming an expert scientist, engineer, or clinician.

This work was supported by the Howard Hughes Medical Institute through an HHMI Professorship grant to C. Wieman.

References

Arum, R. and Roksa, J., (2011) *Academically Adrift: Limited Learning on College Campuses.* Chicago: University of Chicago Press.

Deslauriers, L., Schelew, E., and Wieman, C. (2011). Improved learning in a large-enrollment physics class. *Science*, *332(6031),* 862-864. https://doi.org/10.1126/science.1201783

Ericsson, K. A. (2018). *The differential influence of experience, practice, and deliberate practice on the development of superior individual performance of experts.* In K. A. Ericsson, R. R. Hoffman, A. Kozbelt, & A. M. Williams (Eds.), *Cambridge handbooks in psychology. The Cambridge handbook of expertise and expert performance* (p. 745-769). Cambridge University Press. *https://doi.org/10.1017/9781316480748.038*

Ericsson, K. A., Charness, N., Feltovich, P. J., & Hoffman, R. R. (2006). *The Cambridge handbook of expertise and expert performance.* Cambridge University Press.

Ericsson, K.A., Krampe, R.T., Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, *100(3),* 363-406. https://doi.org/10.1037/0033-295X.100.3.363

Holmes, N.G., Keep, B., & Wieman, C.E. (2020). Developing scientific decision making by structuring and supporting student agency. *Phys. Rev. Phys. Educ, 16(1).* https://doi.org/10.1103/PhysRevPhysEducRes.16.010109

Jones, D.J., Madison, K.W., Wieman, C. (2015). Transforming a fourth year modern optics course using a deliberate practice framework. *Phys. Rev. ST Phys. Educ*, *11(2),* 020108. https://doi.org/10.1103/PhysRevSTPER.11.020108

Lepage, P. (in press). Active learning in a graduate quantum field theory course. *American Journal of Physics.*

Price, A., Kim, C., Burkholder, E., Fritz, A., Wieman, C. (2020). A Detailed Characterization of the Expert Problem-Solving Process in Science and Engineering. arXiv:2005.11463 [physics.ed-ph].

Schwartz, D.L. (2016a). F is for feedback. In D.L. Schwartz, J.M. Tsang, K.P. Blair (Eds.), *The ABCs of How We Learn: 26 Scientifically Proven Approaches, How They Work, and When to Use Them* (pp.64-77). W. W. Norton & Company.

Schwartz, D.L. (2016b). J for just-in-time telling. In D.L. Schwartz, J.M. Tsang, K.P. Blair (Eds.), *The ABCs of How We Learn: 26 Scientifically Proven Approaches, How They Work, and When to Use Them* (pp.114-128). W. W. Norton & Company.

Schwartz, D.L., & Martin, T. (2004) Inventing to prepare for future learning: The hidden efficiency of encouraging original student production in statistics instruction. *Cognition and Instruction, 22(2),* 129-184, DOI: 10.1207/s1532690xci2202_1

Wieman, C. (2019). Expertise in university teaching & the implications for teaching effectiveness, evaluation & training. *Dædalus, MIT Press*, *148(4),* 47-78. https://doi.org/10.1162/daed_a_01760