Chapter 2

Introduction

This chapter discusses the foundations of the research performed for this dissertation. The chapter is organized into four sections. The first section introduces the *LEGO Robot Design Competition*, the workshop class developed as part of this study. The second section discusses the educational goals and motivations that underlay the workshop's development. The third section presents the research methodology upon which the results are based, and the final section introduces the three central chapters of the work.

2.1 A Robot Design Competition

The model class environment that has been developed is in some ways an academic course, a workshop, an independent student activity, and a design competition.[1] The project has taken form as an activity held during MIT's Independent Activities Period (IAP), and is known to students both by its unofficial course number, "6.270," and its name, *The LEGO Robot Design Competition*. The project began as a programming competition modeled after MIT Professor Woodie Flowers' course for Mechanical Engineering students, *Introduction to Design* (course number 2.70). In Professor Flowers' popular course, students each receive a kit of scrap and surplus building materials and the specification for a task to be performed by mechanical contraptions to be designed by the students. Students spend the

[1] This project has been a collaborative effort resulting from the work of many individuals, including, most importantly, Pankaj Oberoi, Randy Sargent, and myself.

two-thirds of the term-long class designing and building machines to operate in the contest held at the end of the class (Flowers, 1987). Flowers' course is unusual in that it gives students wide creative latitude in the execution of their designs. Its high level of popularity amongst the MIT student body can be partly attributed to this feature, which allows the work of the course to adapt itself to each individual student's own learning style.

The format of the LEGO Robot Design Competition project is similar to that of Professor Flowers' mechanical design class. At the beginning of the project, students receive the entire robot-building kit, including plastic LEGO pieces for mechanical and structural work, parts for building a pre-designed microprocessor-based controller that will serve as the "brain" of their robot, electric motors, batteries, and parts to build robotic sensor devices. At the same time as they are given this kit, the specifications for the robot contest are presented. The students' task is to design, assemble, and debug a robot that will successfully accomplish the challenge presented by the contest specification.

The students are given a comprehensive set of course notes that explains how to use the specialized technology they were given in the robot-building kit; also, lectures and recitations are held to present and share ideas relevant to the design and construction of a viable robot. The most important learning, however, happens in small groups and the ever-present laboratory sessions, where students are engaged in building prototypes, testing their ideas, and sharing thoughts with each other and the project organizers.

At the end of the month-long Independent Activity Period, the contest is held. Students robots are on display in a competitive event that is witnessed by about four hundred members of the MIT community. It is an exciting spectacle with an atmosphere resembling that of the traditional collegiate football homecoming game. At the contests, students discover whether their robot can withstand the difficulties of performing in the "real world."

2.2 Relationships for Learning

The development of the Robot Design project as an academic learning environment was guided by a set of pedagogical principles that together form a certain philosophy of education. This approach is founded on the constructionist theory of education espoused by

Seymour Papert (Papert, 1986, 1991). Constructionism is based on the premise that learners are best at creating knowledge and ideas inside their own minds when at the same time they are engaged in an act of building a personally meaningful artifact in the world.

As an elaboration of this principle, I would like the reader to consider a set of educational principles which underlay the Robot Design project as an academic course. The relationships formed by the participants in the Robot Design class—relationships with their work, each other, and the organizers of the project—are different from relationships formed in traditional academic settings. This section highlights these ways in which the Robot Design class differs from traditional academic environments.

2.2.1 Structure and Freedom

In discussion of an educational setting, the word "structure" is often used as a barometer for the degree of guidance in the classroom environment. "Highly structured" situations imply that students' work follows a narrow and precharted path; depending on the teacher's pedagogical beliefs, this may be a good thing or a bad thing. It is usually assumed that such situations produce predictable and consistent results.

On the other end of the spectrum, we hear about "unstructured" learning environments, in which students are free to spend their time as they see fit on an on-going basis. It is generally presumed that the experiences of people in these sorts of environments are much more idiosyncratic than those in the highly structured ones.

Yet the word "structure" is really a misnomer for the sort of variation in different learning situations; one is really concerned with *freedom* and *accountability*. In the "highly structured" situation, the student has little personal freedom and is not terribly accountable for the result, since the learning experience has been precharted by the pedagogue. In the "unstructured" situation—which must necessarily have some sort of structure, however implicit it may be—the participant has a great deal of personal freedom, and also is personally accountable for the result, since the nature of the learning experience is of her or his own making.

Therefore, while the Robot Design project leans far to the side of personal freedom and self-accountability for the learning experience, this is not to say that it is "unstructured." In fact, it is anything but: it has its own consistent internal structure that specifically allows for learning to happen.[2]

In particular, there is the structure of the contest challenge provided to the students. The contest is specifically designed to frame an “problem space” which simultaneously is open-ended yet provides focus to the students.

Open-Ended Problem Spaces

As noted by the MIT Committee on Engineering Design, the generally accepted method of giving students closed-form, “single answer” problems has a significant and detrimental effect on students’ ability to tackle complex and underspecified problems like those in the real world. It is therefore important to give students design problems which do not have a single right answer, but rather a multiplicity of possible good answers.

I prefer to call such situations a *design space* rather than a “problem,” to highlight the open nature of such a situation. A student then does not “solve the contest problem,” but rather constructs a model that satisfies the design space of the contest (to whatever degree of success).

As will be discussed in the Technology for Learning chapter, this premise became a driving concern in the development of the series of contest challenges.

Implicit Structure

While the design space of the contest should be open-ended, it also must set limitations to focus the students’ intellectual activity. Much of the detail in the following chapter explain the way in which the contest specifications, the level of abstraction implicit in the hardware and software, and the nature of the robotics sensors all form a coherent set of constraints and opportunities that guide the students’ work.

[2] I would like to thank Aaron Falbel for his generous conversation with me on this matter.

2.2.2 Accountability

The Robot Design project is designed so that students are on their own to organize their progress through their design project. The only hard milestones are the first day of the course, in which they get their kit, and the last two days of the course, which are the preliminary and final contest rounds. In between this period—i.e., the span of the entire project—students work at their own pace with little formal intervention from the course organizers.

The students are given some suggested milestones, and lectures and recitations presented ideas at what we hoped would be appropriate points along the path of the students' own progress. However, the overarching concept in developing the "course" as it were was to give students all of the materials they needed up front, and get them up and running with the specialized information they might need to know as soon as possible, in order that the control of their progress would be in their own hands. "Get them going and get out of their way" would be a way to summarize the approach. To accomplish this, we analyzed the shortest route between getting the kit of raw parts and getting over the barrier of understanding the basics of a robotic system—sensing, control, and programming. Then we charted a path of progress which would get the students to this point as soon as possible, after which they would be free to manage their own progress.

2.2.3 Versatility

Because of the open-endedness of much of the course, students were free to embark on various sub-projects as they went about building their robots. Some of these became quite significant investigations as they became interested in various aspects of robotics technology. As such, the pedagogical model of the course exhibited a important versatility, adapting to the interests of students from a variety of academic levels and personal backgrounds.

Many of these design excursions will be discussed in other parts of this dissertation: the custom motor drivers built by a number of students in the *Robo-Pong* contest (Section 5.2.1), the burst of musical robots (Section 5.2.3), the custom sensor devices (Section 5.2.2), and the custom control systems (Section 5.2.4). Accomplishing these projects required some mixture of traditional library research, experimentation with the materials that were already provided, and conversations with friends, course instructors, or other experts to get advice.

The framework of the course was to give out a bunch of interesting stuff to play with, along with interesting ideas to explore in doing so, and let the students "go play." If some of them became interested in using a particular idea as a jumping-off point to delve into some associated matter—be it motivated by personal interest, a performance criterion for the robot design, or whatever—then it made the course even stronger if it encouraged them to do just this. Rather than holding students back, in effect saying, "Sorry, that isn't a part of this course, you'll have to do that another time or on your own time," our project said, "Great! Glad to hear you're interested in that. Here are some suggestions. . . " It is a sad fact that many learning situations are so constrained by their predesigned curriculum that there isn't room to allow for individual interests like this.

2.2.4 Teamwork

There were two important ways in which students collaborated with other people during their design work. Not only was their relationship with each other marked by collaboration rather than competition, but so was their relationship to the project organizers characterized more as a peer relationship than as a teacher-student one.

Among Students

Participation in the Robot Design project was different from taking a traditional academic class in the important social dimension. In the Robot Design project, working in a team and living in a fascinated community was a significant part of the experience.

We encouraged students to explicitly think about their strategy for working together. Some teams planned to use a specialist approach, assigning work to each team member based on his or her previous background or expertise. Other teams planned for a generalist approach, in which each team member would share more or less equally in all aspects of the work. Usually teams shifted their strategies as the needs of the project and interests of the teammates became more concrete. For example, a mechanical engineering student might have been elected to perform the LEGO design work, but not long after receiving the kit, all of the team's members realized that they have an interest in developing LEGO designs. Or a team which started out sharing all parts of the work realized that time is running short and that they had better specialize in order to work more efficiently.

We provided students with little supervision in managing their teams, leaving it up to the individuals involved to work out mutually satisfactory arrangements. This resulted in teams that were quite successful, teams that were adequate, and teams that did not hold together. Many students found that getting along with their teammates to be one of the most challenging aspects of the whole project. Students had different strategies for working together on the design project, but for nearly all participants, the challenges and triumphs of working closely with one or two other people became a defining part of their overall experience in the project.

Between Students and Organizers

On numerous occasions, the students of the project developed ideas in partnership with its organizers. This sort of relationship is between teachers and students is not normally highlighted in the academic setting, and had an important impact on the engagement style of the participants. During one of the projects, a student led tutorial sessions to teach students about a favorite sensor of his that was we did not officially support; more recently, the whole project is run by past students who have become teaching assistants and finally organizers.

2.2.5 Community

Beginning in the 1991 contest year, the robot-building software was set up to run on the campus-wide Athena computer network. Since workstation clusters are located not only in various locations on the main academic campus, but also in off-campus living groups (i.e, dormitories and fraternities), students did not need to come to the electronic laboratory and computer cluster to work on their robot projects—they could do so in their living groups.

This happened to some extent in the 1991 contest year, but much more so in 1992. This change is attributable to a difference in the format of the course notes between these two years. In 1991, the course notes were handed out in small segments during the progress of the course. For example, when students first got their kits, they received course notes covering only the contest specification and board assembly instructions. Later, they received handouts on building sensors and using the Interactive C programming language.

Beginning in the 1992 contest year, we were able to provide students with the comprehensive set of course notes from the very start of the class. Since students also received a complete set of electronic assembly tools with their kits, this meant that they could "set up shop" in their living groups and perform a large portion of the robot-building process there, including electronic assembly, LEGO design development, and programming.

As it happened quite a few teams opted to work in this fashion, especially teams that lived off-campus (across the Charles River) or in the west portion of campus (nearly a mile from the electronic laboratory). For these students, the trek to show up at lab wasn't worth the effort if things were going smoothly, especially in January, the dead of Boston's winter.

There was, however, an additional benefit that these teams received from working in their homes: the curiosity, interest, and participation of their housemates. Students described how their housemates would sit down and play with LEGO parts spread over a shared table, peer over their shoulders while they wrote computer code, and get excited watching the fledgling robots wander around the hallways. One year, three teams from one particular dormitory joined forces to build a replica of the contest playing table in their residence for their robot development efforts.

Not only did this type of forum provide students with valuable ideas and feedback, but it gave them a sense that their community cared about the work they were doing. In contrast to most academic endeavors which consist of studying concepts in isolation, this activity was one which they could relish in sharing with their peers.

2.2.6 Motivation

Students' participation in the Robot Design project was purely voluntary. As was mentioned, the project takes place during MIT's Independent Activity Period (IAP), a one-month session between the fall and spring semesters. Student participation in all of IAP opportunities has historically been on an optional basis.

Voluntary Participation

IAP began as an academic experiment in the 1960's, with the purpose of giving MIT students an alternate experience to the traditional university term. IAP has evolved into a veritable smorgasbörd of learning opportunities: faculty, staff, and students organize and run sessions of widely varying formats to introduce both their vocational and avocational interests to others. An IAP activity may consist of a field trip, a lecture series, an all-day workshop, an evening outing, or just about any other format that activity organizers may conceive of and activity participants may find interesting.

It is out of this venue that the Robot Design project was born and bred; presently, the Robot Design activity is by far the largest IAP activity of all. It accepts over one hundred and fifty students, and expects each of them spend between twenty and thirty hours a week on the Robot Design project.[3]

Role of the Contest

Later in this dissertation, the role of the contest in structuring the students' design experiences is discussed. Here, I would like to briefly consider the impact of the contest as a social event.

As a performance event, the contest allows students to display their robots—the fruits of their labor—to the MIT community in a public setting. Because of this, it's natural for students to make an extra effort to make their robots ready for public display, just as anyone puts extra work into preparing his or her work for some sort of performance—be it a concert, talk, or stage play. Students take pride in their creations, and want them to operate properly on the night of the contest, both for their own satisfaction and to impress

[3] I have some qualms about whether the all-consuming nature (in terms of time) of the Robot Design project is appropriate for an IAP experience. As an undergraduate who participated in several IAP's, the smorgasbörd aspect of it was one of the most important for me: that in a month I could partake of perhaps a dozen different learning experiences. Thus I am in part saddened that students who participate in the Robot Design project must necessarily give that up, since the Robot Design work consumes so much of the month's time. I reconcile this conflict by keeping in mind that participation in the Robot Design project is voluntary, so even though students who participate have a different sort of IAP experience, it is still a self-chosen one.

the audience.

Yet a competitive performance is not the only way that a public showing can be staged. It's a natural question to ask if a contest event is an appropriate way to conclude an educational activity; perhaps some students do not feel fully comfortable in an activity that is in some sense evaluated by naming "winners" and consequently "losers." This is an issue that we kept firmly in mind when assessing students' experiences in the class.

We did not want students to become overly serious competitors for two reasons. The first is that this attitude would be most detrimental to those who wished to simply enjoy a friendly, low-key competition. The second is that this would lead students to work in isolation rather than in collaboration with each, and would have a negative effect on their learning potential.

In light of this, we took several actions to encourage friendly rather than ruthless competition. We discontinued a policy of having substantial prizes (e.g., VCRs, scientific calculators, and cordless telephones) for the contest winners that was inherited from the Robot Design project in its first years. We established a rule stating that robots could not intentionally harm one another, encouraging the competitions to be more like sporting events based on finesse and speed (e.g., tennis) than brute strength and machismo (e.g., American football). We encouraged students to take a light-hearted approach to the competition and share ideas rather than working in isolation; in our experience, we told them, teams that were present in the lab, sharing ideas and discussing problems, were more likely to have a working and reliable robot than those that worked on their own, hoping to develop a "world beater" strategy that would catch the other designs off-guard.

We were mostly successful in creating an atmosphere of friendly sport. While there were always a few teams who put winning ahead of most other aspects of participation, generally speaking this was not the case. While some teams developed their robots in isolation, in an attempt to surprise their competition, most students worked in a public setting. Others developed robots that were "accessorized" in such a way that would be hard to imagine that winning were the primary concern—for these students, creating an interesting project with a design aesthetic meaningful to them was the overriding concern.

An important feature of the contests was the fact that they provided an objective performance evaluation for the students to base an analysis of their work. This is in contrast to events like the talent show, where the rating of a panel of judges is the evaluation method; this can hardly be considered objective in the scientific sense. For our Robot Design contests, however, the judges were more like game officials whose task it was simply to assist in procedural matters of running the game and ascertaining that the rules were adhered to. For the cadre of MIT students involved, it became apparent that this sort of game had great appeal—both from a standpoint of participation and observation.

Because the contest required an *actual performance*, not simply a theoretical or proof-of-concept one (e.g., “There; it worked once, now I’m done”), interesting issues were forced to the surface during the design process. Problems like reliability, endurance, and overall robustness are things that students would not have faced, by and large, in developing a project that did not require both a public demonstration and one that had hard-and-fast performance goals. Further, it was these issues in particular that became central places of learning for many of the project participants: demonstrating that something worked “in principle” was a familiar task, but proving that it worked in repeated practice was something altogether different. In this manner, the project gave them experience comparable to design activities in the real world.

The motivational aspect of knowing that one has an excited audience to receive one’s work should not be underestimated. In recent years of the course, all of the project participants knew how popular the contest events were, and there is no question in my mind that this served as an inspiration for doing one’s best. Even when students might realize that their machine is not likely to be the overall contest winner, students still have the desire to get the machine to work as they intended it to, and thereby be a winner in their own minds.

2.3 Methodology and Evaluation

The methodology used in this work is based on an iterative cycle of design, testing, and evaluation, as represented in Figure 2-1. In the first stage, a hypothesis about the essential characteristics of a design-rich learning environment is formed. As was mentioned, this hypothesis was based on Papert’s theory of constructionist learning.

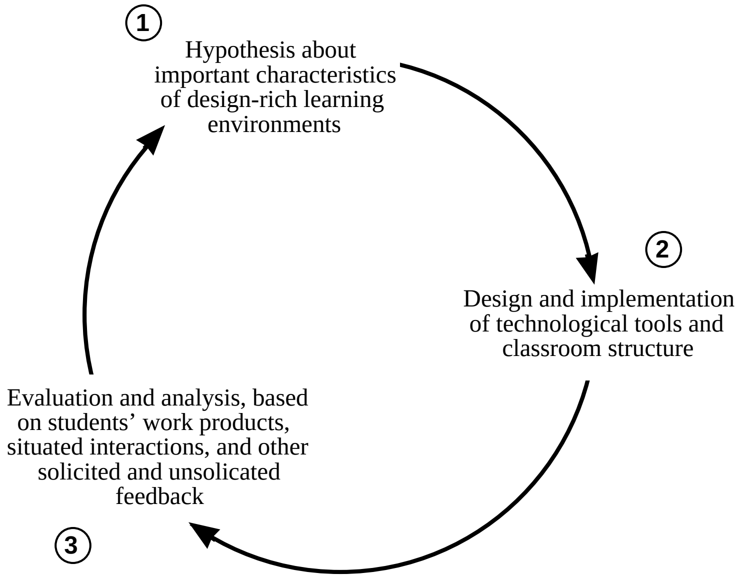

Figure 2-1: Cycle of project design and assessment

In the second stage, this hypothesis guides the formation of both a set of technological tools for learning and the classroom environment in which they are deployed. In this work, technology to facilitate exploration in the areas of engineering, design, sensing, and control was developed.

In the third stage, these materials are tested, evaluated, and analyzed. A variety of observational tools are used to form the evaluation: situated interactions with students during the course of their design projects, student journals and other solicited comments, unsolicited comments, video reports, and interviews, and analyses of the students' actual work products—robotic designs that reveal structural, software, and strategic decisions.

We have completed four iterations of this cycle of design and evaluation, corresponding to the four successive versions of the Robot Design course which were developed as part of this dissertation. In each of these versions, the impact of the materials provided to the students upon their learning was assessed, and these analyses fed into the design of the materials for the subsequent course.

The question of how to accurately assess what a person has "learned" through the progress of a particular experience is a difficult one. Traditional education bases its evaluation upon students' performance on a series of written tests, which are usually either of the "take-home" (i.e., problem set) or in-class (i.e., examination) variety. If a student is proficient in his or her performance on these measures, then it is assumed that he or she has learned the material that is the object of the class.

This model is problematic from a constructionist learning theory point of view. The written examination is only one of what could be numerous "output modalities" for a student's knowledge; in the traditional model, this mode is so heavily weighted in value so as to exclude others. Put another way, there is not a straightforward mapping between what a student knows and what he or she is able to express on a written test.

Further, this focus on written evaluation throughout the educational system forces the implication that all important knowledge or skills can be represented in this form. This in turn implies that any knowledge which cannot be represented this way is not worth knowing (at best, it is not worth evaluating). Yet it is precisely these alternative skills and ways of knowing—what Schön refers to as the tacit knowledge embedded in a designer's process of work—which are the locus of investigation here. In other words, I am less concerned with what students can write or say about what they have learned than with a textured understanding of the difficulties they have faced, the styles of work they established, and the problems they have posed and solved for themselves.

These assessments of students' activity are based on data collected from several means:

Situated interactions with students. During the course of their projects, I interacted with students extensively, answering their questions, helping them solve problems, and otherwise listening to their thoughts and concerns. This data was essential for a detailed understanding of the thought processes of individual students. As noted earlier, this style of interaction was more as between peers than as between teacher and student; often the problem puzzling the students was one we had not yet solved ourselves.

Student journals and weekly video reports. Weekly journals and self-guided video reports kept by students recorded the progress of their designs. This data was useful in understanding the broadness and commonality of students' experiences.

Concluding video interviews. Most teams were interviewed near the end of the course in a relaxed setting. This data was useful for recording students' own descriptions of their experience (which did not necessarily match with other data used to observe it).

Analysis of student projects. Records kept on the students' designs—including copies of their control software, photographs of their machines, and discussions of the designs in the journal reports—is analyzed to infer the sorts of problem situations the students framed for themselves during the course of their designs, and how these challenges were resolved.

This data is used as the basis for forming arguments about learning, but also serves as feedback into the cyclical process of course design mentioned earlier. In this regard, the inferences made from these observations are either strengthened or challenged by seeing the results obtained from applying the working results into the course design. Thus, the final arguments to be presented in this work are based on their value as tested in the course itself.

It is worth mentioning that students often cannot verbally explain their learning process. In the concluding interviews, there were a number of students who were quite animated and involved during the progress of the course—excited about the new ideas they were experiencing and learning. But when it came to the interview, they were at a loss for words, unable to describe what had been valuable or interesting to them. Similarly, many students found it much easier to discuss the strategy their robot used to solve the contest than the strategy *they* had employed to create this robot.

For this reason I cannot overemphasize the importance of the contextualized interactions I had with students in my role as a course organizer. It was when working with students on problems they were dealing with at the particular moment that I gained the greatest understanding of the issues they faced and how they shaped the problem situation for themselves.

A related matter is the fact that sometimes learning happens over a long period of time. In one of the case studies discussed in Chapter 5, a student's work is tracked over a period that spans about fifteen months. It is only at the end of the this period that the student in question revises his fundamental beliefs, after a series of demonstrated failures. The message of this example is that students' ideas are deeply embedded and it takes time for them to change; sometimes the failure of one robot is not proof enough.

2.4 The Main Themes

The section introduces three themes that are the subjects of the subsequent three chapters.

2.4.1 Technology for Learning

Chapter 3, entitled “Technology for Learning,” explores the development and impact of the technology created for the Robot Design project. Taken in a broad sense, “technology” includes not only the tangible hardware and software materials that were given to the students, but the specifications of the robotic contest challenges. The chapter explores both the types of problems students encountered in the their robot-building activity, and the progressive way in which we gained a deeper understanding of the structure that underlay the design environment of the robot-building task.

Several themes emerge from this study. The first of these relates to the way in which the contest game specification affected students’ work. We found that it was difficult to balance our desire for an open contest specification, which would allow many different types of robotic solutions, which a focused specification, which would encourage students to make machines that really worked.

Another issue relates to the properties of the technology as an educational media. Specifically, characteristics like *interactivity*, *modularity*, and *transparency* were found to have a profound impact upon students’ learning.

A third theme arises from unexpected difficulties in effectively using sensor devices. In two cases, complications in using sensors put the project participants and the organizers together in collaborating to find a way to use the devices. This “leveling of the playing field” had a valuable influence on the relationships of all of the project participants.

2.4.2 Ideal and Real Systems

In Chapter 4, entitled "Ideal and Real Systems," the students' robots are analyzed from a control systems point of view. Of particular interest is students' recurring inclination to build robots that will perform properly only under ideal conditions. Students repeatedly build robots that are not well-equipped to deal with the exigencies of the real world, but rather with the specifications of an idealized, abstractified world, a world that they would like to believe is a close representation of reality. This result points to limitations in the set of ideas about technological systems and methods that comprise the core of the engineering curriculum; what surprises many participants is that these ideas do not map well to the challenge of designing a robot to play in one of our contests.

2.4.3 Design Styles

In Chapter 5, entitled "Design Styles," I explain how the course attempted to encourage a bottom-up, playful style of design—the "bricoleur" style described by Turkle and Papert. We found that many students gravitated toward this style in the mechanical design aspect of their robotics work, but returned to the apparently more familiar top-down style, which one student described as the "computer science mind."

This chapter documents a variety of design excursions, mini-research projects students performed as offshoots of their robot-building work, and ways that students re-defined the goals of their participation to suit their own interests. These projects illustrate the flexibility of the project in stimulating genuine participation from its students.

The data from this chapter is used to argue that regardless of students' approach toward design, they encounter the central engineering design problem of reconciling the phenomena that they observe with the formal models of engineering science that they have learned. In some cases, it is easy for them to see the underlying principles, but in many it is not.