2013

Mathematical Knowledge for Teaching and Visualizing Differential Geometry

Nathan Pinsky
Harvey Mudd College

Mathematical Knowledge for Teaching and Visualizing Differential Geometry

Nathan Pinsky

Weiqing Gu, Advisor

Darryl Yong, Reader

HARVEY MUDD COLLEGE

Department of Mathematics

May, 2013

Abstract

In recent decades, education researchers have recognized the need for teachers to have a nuanced content knowledge in addition to pedagogical knowledge, but very little research was conducted into what this knowledge would entail. Beginning in 2008, math education researchers began to develop a theoretical framework for the mathematical knowledge needed for teaching, but their work focused primarily on elementary schools. I will present an analysis of the mathematical knowledge needed for teaching about the regular curves and surfaces, two important concepts in differential geometry which generalize to the advanced notion of a manifold, both in a college classroom and in an on-line format. I will also comment on the philosophical and political questions that arise in this analysis.

Contents

List of Figures

Acknowledgments

First and foremost, I would like to thank my thesis advisor, Prof. Weiqing Gu, for her guidance and support on this project, and my second reader, Prof. Darryl Yong, who met with me several times to discuss ideas in mathematics education and kept me pointed in the right direction. I would also like to thank Claire Connelly, who on numerous occasions helped me with technical problems and got my LaTeX to compile, and Prof. Talithia Williams, for her moral support.

Chapter 1

Introduction

This project is motivated by interest along two parallel lines of inquiry: pedagogy in distance education courses and pedagogy in higher education. Distance education has existed for quite a while in the form of correspondence courses, but as the Internet has matured, so too has the prominence of such courses. Perhaps the best known example of Internet distance education is the Khan Academy, which offers a host of free video tutorials in a vast range of subjects for its users' self-edification. However, a number of universities are experimenting with Internet-based learning in a more rigorous framework, the massive open online course (MOOC). MOOCs, which are offered through companies run by such universities as Harvard, MIT, and Stanford, are free and have no enrollment limit, meaning that elite college-level education is more accessible to people worldwide (Lewin, 2012). Although MOOCs have many benefits, including worldwide feedback on the ideas of both the instructor and the students, opinion is divided on whether any type of credential should be awarded for completion of the course. One reason for this controversy is that this type of distance education is still experimental. Very little research has been done in distance education pedagogy; in fact, little research has been done regarding pedagogy in higher education at all. This project intends to contribute to the effort to understand such pedagogy by examining teacher knowledge in higher education and distance learning in the context of a mathematics course.

Today it is widely accepted that the knowledge necessary for teaching a subject goes beyond that necessary for using or studying the same material, even in a research setting. A teacher cannot be limited to simply knowing the principles, values, and main results of a field, but must also understand why a result is valued or considered valid, how it fits into the historical arc of a particular field, and what might be considered a legitimate objection to such a result, and be able to articulate all of these ideas to students who have considerably less academic experience. This conceptualization of the richness and depth of teacher knowledge was first put forth by Shulman (1986) in an address, later published, in which he argued for a study of teachers' content knowledge to guide educational policy and teacher preparation programs. His idea, termed *pedagogical content knowledge* (PCK), has since been developed by many researchers and, in the context of mathematics, encompassed into what is known as *mathematical knowledge for teaching* (MKT). As the name suggests, MKT refers to knowledge of mathematics that is needed to teach it well. MKT is currently recognized as containing both subject matter knowledge, including *specialized content knowledge* (SCK), and pedagogical content knowledge, which includes *knowledge of content and students* (KCS) and *knowledge of content and teaching* (KCT); for a thorough explanation of these constructs, see Chapter 2.

While MKT has proven to be a valuable conceptual framework in analyzing teacher knowledge (see, e.g., Beswick, 2012; Charalambous and Hill, 2012; McCrory et al., 2012), most research has been done at the elementary level and in classrooms (see, e.g., McCrory et al., 2012; Speer et al., 2010). Speer and Hald (2008), some of the few researchers interested in MKT at the undergraduate level, provide an overview of this research, and conclude,

> To date, analogous studies have not been conducted with people who teach undergraduate mathematics. It remains to be seen what kinds of specialized knowledge of mathematics graduate students and professors develop as a result of interpreting students' ideas and engaging in other teaching-related mathematical activities (Speer and Hald, 2008: p. 308).

Simply put, we do not know what MKT looks like at the undergraduate level. Undergraduate students are significantly more mature than elementary, and sometimes even secondary, students both emotionally and mathematically, and it is as yet unknown how this mathematical maturity affects required instructor knowledge. Examples of MKT from both elementary and post-secondary education are given in Chapter 2, but these examples are limited; Speer and Hald, from whom the post-secondary examples are drawn, focus on "the mathematics taught in colleges (e.g., college algebra, pre-calculus calculus) [that] is also taught in high school" (p. 304). There is, however, a great deal of mathematics taught at the undergraduate level beyond that taught in high school, adding yet another gap in mathematical maturity. In this project, we ask the question, "What does a teacher need to know about an upper-division mathematical topic to teach it well to mathematically mature students?" We can even take this pursuit one step further, asking how such knowledge would change if the teacher had limited interaction with students, such as in a large lecture class or in a MOOC. Such research could help professors plan better lessons, increase participation in and enthusiasm for higher mathematics, and shed light on teaching to advanced students below the undergraduate level.

To address these questions, we examine the teaching of an introductory course in differential geometry. Differential geometry lies at the intersection of many different areas of mathematics, including calculus, linear algebra, analysis, topology, and differential equations. As a result, only students with significantly strong mathematical backgrounds take this course, and the teaching of it draws on teaching from many different areas. It is, however, more than the sum of these parts—differential geometry requires an understanding of space that is frequently very new even for students well versed in all these areas. Successful work in differential geometry also relies heavily on visualization of the mathematical ideas, which can be facilitated by mathematical software packages like Mathematica. It is here that our interest in distance learning comes back into play; the examples generated in a Mathematica-based differential geometry course would lend themselves naturally to a video or demonstration that could be put online for widespread use. The generation of such lessons would therefore need to account for students learning the material in a distance-education setting.

However, an introductory course in differential geometry differs from other courses not just in the level of the content presented but also in the context in which that content is presented. Compared to classes in elementary, middle, or high school, where attendance is compulsory and a curriculum is mandated by state or federal law, the differential geometry class we consider here is an upper-division elective at a postsecondary institution. Thus, not only is enrollment in the institution a privilege, but attendance in the class is based on individual student interest. Moreover, because the curriculum of the course is not governed by any external regulations, both the content covered and the guiding goals of the course are at the instructor's discretion. In such an advanced class, one of the primary goals of the course is to prepare students for graduate-level study, including self-guided investigation and individual effort to fill in gaps in the content or reasoning presented in the course. This goal disappears, of course, in a distance-education context, providing us with a comparison with which we may evaluate the relative effects of the content and context on MKT.

In order to investigate the MKT present in this undergraduate differential geometry course, we present a thorough review of existing research on MKT in the classroom in Chapter 2. We start by presenting the historical and theoretical origins of pedagogical content knowledge and how it has been developed into the framework of MKT used today. We then more closely examine this framework, with a particular emphasis on how it is understood and applied by researchers. Numerous examples from recent research allow us to more completely define each subcomponent of MKT. In Chapter 3, we present our analysis of introductory differential geometry within the theoretical framework developed in Chapter 2. This analysis is largely theoretical, but is also supported by examples from lessons on regular curves and regular surfaces for a distance education-style course. Finally, in Chapter 4 we comment on the pedagogical and philosophical implications of our findings for both teaching and further research.

Chapter 2

Theoretical Background: Pedagogical Content Knowledge and Mathematical Knowledge for Teaching

2.1 Pedagogical Content Knowledge

Pedagogical content knowledge (PCK) was first proposed by Lee Shulman in a 1985 Presidential Address at the annual meeting of the American Educational Research Association, later published as a paper in *Educational Researcher* in 1986. Pedagogical content knowledge as an idea caught on like wildfire—according to Ball et al. (2008), Shulman's two articles on the subject (1986; 1987) have over 1200 citations in refereed journals including at least 50 per year between 1990 and 2008 in many different subject areas, and in recent years has been the backbone of a theoretical framework for research into what teachers know. Shulman's original address, however, was a response to, and commentary on, the educational reforms of the 1980s. According to Shulman, although many states required examinations for teacher licensure, the content exams were little more than "tests of basic abilities to read, write, spell, calculate, and solve arithmetic problems. . . in most states, however, the evaluation of teachers emphasize[d] the assessment of capacity to teach" (Shulman, 1986: p. 5). Although it semeed standard in the 1980s, this lack of content examination was not always the case. Shulman explains that, in 1875, elementary teachers were required to take an examination covering 20 different topics, and the "Theory and Practice of Teaching" (that is, pedagogy) accounted for only 50 of the possible 1000 points (Shulman, 1986: p. 4-5). However, the focus on pedagogy in the 1980s was not constrained to teacher assessment policy, but extended to the research community as well. In reality, the research community's focus on pedagogy was likely the cause of the policymakers' decisions, as the most recent research, on which policymakers based their standards, was focused entirely on pedagogy.

Shulman's goals were to even out this imbalance by generating a renewed emphasis in the research community on teacher knowledge of the subject matter they teach. Shulman and his colleagues in the "Knowledge Growth in Teaching" research program studied young secondary teachers in English, biology, mathematics, and social studies as they developed as teachers, aiming to "trace their intellectual biography—that set of understandings, conceptions, and orientations that constitutes the source of their comprehension of the subjects they teach" (Shulman, 1986: p. 8). Instead of applying standardized tests to measure this knowledge, however, Shulman took novel approach guided by the following revolutionary questions (Shulman, 1986):

- What are the sources of teacher knowledge?
- What does a teacher know and when did he or she come to know it?
- How do teachers decide what to teach, how to represent it, how to question students about it and how to deal with problems of misunderstanding?
- In the face of poor resources (e.g., a confusing textbook), how does the teacher employ content expertise to generate new explanations, representations, or clarifications?
- How does the teacher prepare to teach something never previously learned (or learned long ago)? How does learning *for* teaching occur?
- How do teachers take a piece of text and transform their understanding of it into instruction that their students can comprehend?

Although his research was ongoing at the time of publication, Shulman suggested three categories of content knowledge: subject matter content knowledge, curricular content knowledge, and pedagogical content knowledge. Of these three, pedagogical content knowledge has spurred the most interesting research in teaching, and, as the basis for mathematical knowledge for teaching, it is most relevant to the work in this thesis. As with any soft science, these categories are not entirely well-defined, and there may be overlap between them. Nevertheless, we feel that it is important to delineate all of these categories in order to make as clear as possible what pedagogical content knowledge does and does not encompass.

According to Shulman, *subject matter content knowledge* "refers to the amount and organization of knowledge per se in the mind of the teacher." However, the teacher needs to know more about the subject matter than the content itself. To teach well, one must know "the variety of ways in which the basic concepts and principles of the discipline are incorporated to incorporate its facts," and "must be able to explain why a particular proposition is deemed warranted, why it is worth knowing, how it relates to other propositions, both within the discipline and without, both in theory and in practice" (Shulman, 1986: p. 9). That is, in order to teach a subject, one must understand it on a much deeper level than is needed simply to use it. One must understand the strength of, the reasoning behind, and the relative importance of common principles that the professional appeals to daily without reflection. Although this may not be surprising today (Einstein, for example, is frequently credited as saying that "you do not really understand something unless you can explain it to your grandmother"), Shulman takes great pains to establish this in his philosophical and historically-based conceptualization of teachers as "masters" or "doctors." Importantly, though, subject matter content knowledge is knowledge of the subject itself. In this regard, it is the most purely academic of Shulman's categories.

By contrast, *curricular content knowledge* is the category of content knowledge most situated in practice. Shulman describes the curriculum as

> the full range of programs designed for the teaching of particular subjects and topics at a given level, the variety of instructional materials available in relation to those programs, and the set of characteristics that serve as both the indications and contraindications for the use of particular curriculum or program materials in particular circumstances. (Shulman, 1986: p. 10)

In this regard, curricular content knowledge is primarily a knowledge of the large scale tools that can be used to teach content. These might be established syllabi and texts, both mainstream and alternative, films, or classroom demonstrations. However, the curriculum is not limited solely to the content in one particular course. Shulman also specifies two additional aspects of curricular content knowledge, *lateral curriculum knowledge* and *vertical curriculum knowledge*. Lateral curriculum knowledge refers to knowledge of other classes the student is currently taking, while vertical curriculum knowledge refers to knowledge of past and future classes in the same subject area. In both cases, this knowledge allows a teacher to build connections between concepts and prepare students for future studies.

The last and, for the purposes of this thesis, most important category of knowledge proposed by Shulman is *pedagogical content knowledge*. Shulman explains pedagogical content knowledge (PCK) as going "beyond knowledge of subject matter per se to the dimension of subject matter *for teaching*...the particular form of content knowledge that embodies the aspects of content most germane to its teachability" (Shulman, 1986: p. 9, emphasis in original). This, according to Shulman, includes "the most useful forms of representation...in a word, the ways of representing and formulating the subject that make it comprehensible to others" as well as "an understanding of what makes the learning of specific topics easy or difficult: the conceptions and preconceptions that students of different ages and backgrounds bring with them to the learning" (Shulman, 1986: p. 9). Shulman suggests that, especially for common topics, a teacher's pedagogical content knowledge would include "the most powerful analogies, illustrations, examples, explanations, and demonstrations," but acknowledges that none of these will ever be all-powerful or ubiquitous, and reasons that teachers must therefore "have at hand a veritable armamentarium of alternative forms of representation" which arise both from research and from "the wisdom of practice" (Shulman, 1986: p. 9). He also recognizes that students come to any lesson with a priori knowledge, and that a teacher must know how to use these ideas to shape a lesson. Shulman's conception of pedagogical content knowledge, therefore, is based as much in practice and understanding of student learning as it is in research and understanding of teaching methods.

2.2 Mathematical Knowledge for Teaching

Shulman's work fundamentally changed how many researchers thought about teaching. As noted above, Shulman's two papers have been cited over 1200 times since their publication; he effectively established a new paradigm of educational theory. However, the working definitions of pedagogical content knowledge in research that built on Shulman's were as broad and unclear as his original explanation until it was refined and operationalized by Ball et al. (2008). Specifically, Ball et al. note that, prior to their work, most working definitions of PCK were "broad enough to include nearly any package of teacher knowledge and beliefs" (Ball et al., 2008: p. 394) and had insufficient empirical backing; the resulting work therefore was entirely hypothetical and could not be used effectively to guide policy or teacher preparation (Ball et al., 2008: p. 390).

In what is now a cornerstone paper, Ball et al. (2008) begin to address this problem by examining "what teaching itself demands" (p. 390). This emphasis on actual teaching, as opposed to curriuculum, is a direct response to the issues identified above. Curriculum is based on existing ideas about the relative importance of pedagogy and content knowledge, and is therefore guided by the same untested hypotheses, developed through "logical and ad hoc arguments about the content believed to be necessary for teachers" (Ball et al., 2008: p. 390), which these researchers were investigating. It is fundamentally normative, rather than empirical. Instead, they looked at what actually goes on in classroom lessons, guided by the question "What do teachers need to know and be able to do in order to teach effectively?" They explain,

> This places the emphasis on the use of knowledge in and for teaching rather than on teachers themselves... In other words, although we examine particular teachers and students at given moments in time, our focus is on what this actual instruction suggests for a detailed job description (Ball et al., 2008).

Simply put, Ball et al. "lay the foundation for a practice-based theory of mathematical knowledge for teaching" (p. 395).

The empirical research on which Ball et al. built their theory was conducted in two phases. In the first, they conducted "extensive qualitative analyses of teaching practice" (p. 395). They examined collections of classroom documentation, most notably a "large longitudinal National Science Foundation-funded database, documenting an entire year of the mathematics teaching in a third grade public school classroom during 1989-1990," which included video and audio recordings, transcripts, copies of individual student work and teacher notes, and curriculum materials. This work did not focus solely on single classroom sessions but on the progression of the teaching and learning through a full year. The researchers also drew on "the wide range of experiences and disciplinary backgrounds of the the members of [their] research group" and a "set of analytic tools... developed for coordinating mathematical and pedagogical perspectives." Unfortunately, the article provides only an unpublished manuscript as reference for these tools. In the second phase of their research, in coordination with the Study of Instructional Improvement (www.sii.soe.umich.edu), Ball et al.

"began to develop and validate survey measures of mathematical knowledge for teaching" (p. 396). The most important result of this work was the confirmation of their belief that MKT is "multidimensional. That is, general mathematical ability does not fully account for the knowledge and skills entailed in teaching mathematics" (Ball et al., 2008).

Mathematical knowledge for teaching, in Ball et al.'s formulation, is in fact incredibly complex. In order to illustrate it, they present as an example the analysis of student work, including errors, on subtraction of three-digit numbers. They show a number of different errors, with the minimal annotation that one would expect on student work, and explain how a teacher must recognize wrong answers, recognize common procedural errors, recognize and process unfamiliar processes, and speak fluently to each one's validity, all in real time in front of a room of students. In their words,

> Teachers must know rationales for procedures, meanings for terms, explanations for concepts. Teachers need effective ways of representing the meaning of the subtraction algorithm—not just to confirm the answer but to show what the steps of the procedure mean and why they make sense... Teaching also involves considering what numbers are strategic to use in an example. The numbers 307 and 168 may not be ideal choices to make visible the conceptual structure of the [subtraction] algorithm. Should the numerical examples require two regroupings, as in this case, or should examples be sequenced from ones requiring no regrouping to ones that require several? What about the role of zeros at different points in the procedure? Should the example include zeros—or perhaps not at first? (Ball et al., 2008: p. 398).

Although these types of thoughts do not usually come to mind when one thinks about subtraction (and especially not every time one subtracts!), they are each vitally important to consider when teaching beginning mathematicians the subtraction algorithm. To more effectively categorize the types of thinking and knowledge that constitute MKT, Ball et al. introduce several subdomains, shown in Figure 2.1, which lie on a continuum from subject matter knowledge to pedagogical content knowledge. Common content knowledge (CCK) refers to a knowledge of the subject matter that is used in many settings; in the subtraction example given above, common content knowledge would refer to simply knowing how to subtract. We hope that the importance of this type of knowledge is self-evident, so

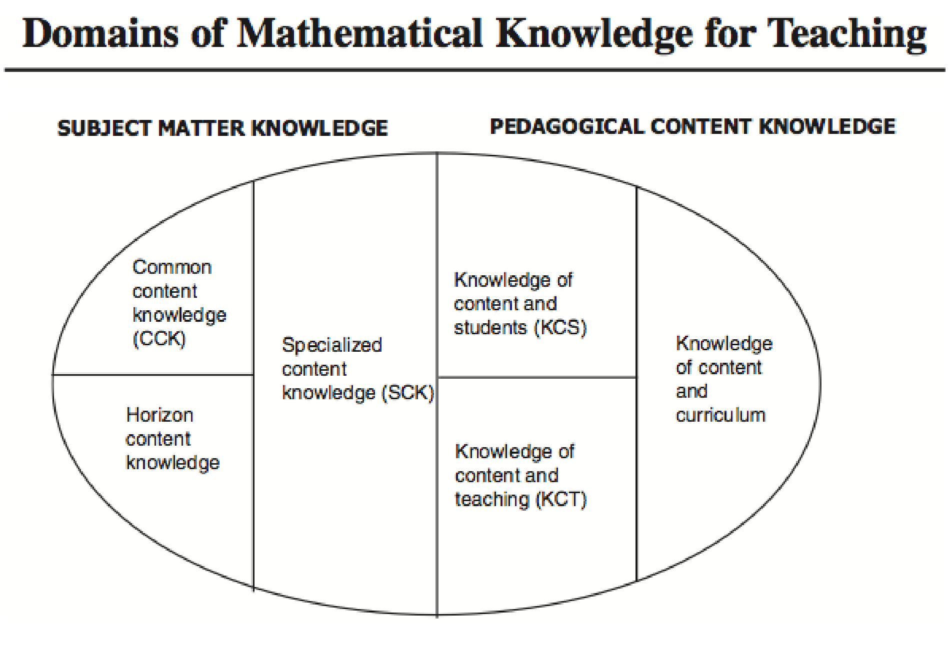

Figure 2.1 Graphical summary of Ball et al. (2008)'s subdomains of MKT, taken from that paper (p. 403).

we do not go into more detail here. We also note that two of the categories, *horizon content knowledge* and *knowledge of content and curriculum*, are included as separate categories only provisionally. Horizon content knowledge is "an awareness of how mathematical topics are related over the span of mathematics included in the curriculum" (Ball et al., 2008: p. 403), which we might expect to come naturally with common content knowledge of later topics, while knowledge of content and curriculum is simply a renaming of Shulman's curricular content knowledge. Ball et al. remark that both of these categories "may run across several categories or be a category in [their] own right," and express their "hope to explore these ideas. . . as [they] are used in teacher education or in the development of curriculum materials for use in professional development" (Ball et al., 2008: p. 403). The other three categories, *specialized content knowledge*, *knowledge of content and students*, and *knowledge of content and teaching*, however, are much more clearly defined, and will play important roles in our analysis.

Specialized content knowledge (SCK), which refers to "mathematical knowledge and skill unique to teaching" (p. 400) is perhaps the most interesting result from Ball et al.. They found that, in the process of teaching, teachers must access mathematical knowledge in a "decompressed" form in order to accomplish tasks specific to teaching, such as responding to students' "why" questions, linking representations to each other and to underlying mathematical concepts, or using mathematical notation and language and critiquing its use. While the goal of teaching is to enable students to use mathematical constructs comfortably and apply mathematical strategies confidently without worrying overtly about why they work, the teacher must worry about such things so that they can "mak[e] features of particular content visible to and learnable by students" (Ball et al., 2008). For example, while "engineers have to mathematically model properties of materials," they need not "explain why, when you multiply by 10, you 'add a zero' " or "identify and distinguish the complete range of different situations modeled by $38 \div 4$" (Ball et al., 2008: p. 401). This type of knowledge in fact resembles that used by mathematicians engaged in research in pure mathematics. In both cases, one is concerned with finding appropriate definitions for abstract concepts, and must frequently contort their own mental representations in order to accomplish their goal. While unpacked and prompted by tasks required of teachers, SCK is still distinctly knowledge of the subject matter, and therefore is positioned under subject matter knowledge, but closer to pedagogical content knowledge.

Knowledge of content and students (KCS), on the other hand, lies on the subject-matter-knowledge end of pedagogical content knowledge. This type of knowledge, as the name suggests, relates to an understanding of students and how they will interact with the content. For example, a good teacher will try to anticipate which sections of a lesson will be easy for the class to understand, and why they might struggle with another, and will try to choose an example that will interest students (Ball et al., 2008: p. 401). However, this type of knowledge is also used when listening to students or reading their work. Teachers must be able to process student ideas, even when expressed in unexpected or incorrect forms, and analyze them for understanding or lack thereof; they must be able to recognize mathematical ideas even in very lax contexts. Ball et al. (2008) illustrate the difference between CCK, SCK, and KCS:

> Recognizing a wrong answer is common content knowledge (CCK), whereas sizing up the nature of an error, especially an unfamiliar error, typically requires nimbleness in thinking about numbers, attention to patterns, and flexible thinking about meaning in ways that are distinctive of specialized content knowl-

> edge (SCK). In contrast, familiarity with common errors and deciding which of several errors students are most likely to make are examples of knowledge of content and students (KCS) (p. 401).

The last subdomain identified by Ball et al. is knowledge of content and teaching (KCT), an applied knowledge of pedagogical principles to the content material. It is the knowledge required for sequencing topics for instruction, evaluating multiple representations for presentation, and understanding the benefits of presenting one, some, or all of them. For example, a teacher needs to know the relative merits of using coins, counting beans, and base-ten blocks for teaching addition and subtraction, and be able to use each in different situations for specific pedagogical purposes. More generally, "each model. . . requires different care in use in order to make the mathematical issues salient and useable by students. . . Knowing how these differences matter for the development of the topic is part of [KCT]" (Ball et al., 2008: p. 402).

These domains of mathematical knowledge for teaching are, of course, not precisely delineated; Ball et al. acknowledge that "it is not always easy to discern where one of our categories divides from the next, and this affects the precision (or lack thereof) of our definitions" (p. 402). Any given task could quite easily require knowledge from many or all of these domains. The importance of Ball et al.'s work resides not in clearly drawing the boundaries between them, but rather in establishing their existence. For establishing useful boundaries, we turn to examples from research that has built on this work.

2.3 Applying The Theory

To further clarify what is meant by each of SCK, KCS, and KCT, we consider an example of recent research on the subject. We note, however, that these and many other authors invent their own vocabulary based on that of Ball et al., and so some liberties must be taken to pinpoint which dimension of MKT was meant. As the study of mathematical knowledge for teaching is relatively new, this is not altogether surprising.

2.3.1 Speer and Wagner (2009)

Speer and Wagner (2009) analyze how MKT (or the lack thereof) affects a professor's ability to effectively run a discussion-based introductory course in differential equations at the undergraduate level. They chose this context because it allowed them to study "teachers whose mathematical content knowledge is very strong, thus allowing a clearer picture of what other types of knowledge teachers at all levels may need in order to adopt new teaching practices aligned with contemporary reform curricula" (p. 533-4). Notably for this project, the course also used "Java applets designed to guide students through discovery of the core concepts of a dynamical systems approach to Differential Equations" (p. 538), allowing us to examine the effect of computer-based lessons on student learning, and therefore the things that a teacher must consider in designing such a program. Their analysis, like ours, stems largely from examinations of lessons after delivery and hypothetical extrapolations of how certain types of knowledge might have changed the events. They present four episodes from this class, organized into two case studies.

In the first case study, titled "When Students Contribute Good Ideas," Speer and Wagner consider episodes in which the professor reacts to mathematically productive ideas. In the first episode, students are discussing whether, if the rate of change of a population is dependent only on the population size and the population size is dependent on time, they can say that the rate of change of the population is dependent on time (that is, if $dP/dt = f(P)$ and $P = P(t)$, then is it true that $dP/dt = g(t)$?). The authors assert that the professor's CCK caused him to focus on the mention of initial conditions, but "did not have the PCK to enable him to recognize that the issues were relevant to the development of students' understanding" (p. 544). Although they do not explicitly specify what type of PCK they are referring to, it can be inferred from their repeated references to "students' understanding" that this is KCS. We also note that the symbolic unpacking of the discussion given above, as well as the ability to frame the students' problem as such, is a prime example of SCK.

The second episode in this case study, in which students are discussing the difference between the differential equations $dP/dt = P$ and $dP/dt = e^t$, provides a second example of SCK. In this episode, a student makes a comment that, although confused, reveals understanding of the material, but the professor does not recognize it. Speer and Wagner note that "it is SCK that supports the recognition of valuable student contributions without having to do so much mathematical 'work' in the moment while so many other issues are simultaneously demanding attention" (p. 548). They then go on to assert that, had the professor been able to suggest $P(t) = 2e^t$ as a function to examine, the conversation might have been more productive; they identify this as the missing SCK.

In the second case study, titled "When Students Pose Not-So-Good Ideas," Speer and Wagner consider episodes in which the professor cannot steer unproductive conversations into productive discussions. In the first episode, students examine first-order linear differential equations generated to suggest an application of the product rule, but fail to see the connection. In the second episode, in which the class is considering the same equations, some students suggest the use of integration by parts rather than the chain rule. They suggest that, in both episodes of this case study, the professor "lacks the SCK to follow the students' ideas and the PCK to envision how their ideas, even though they were incorrect, could have been used productively" (p. 550). We also note that, based on the transcripts, the example equations did not effectively suggest an application of the chain rule. Although the professor did write the equation in two forms to direct students' attention to particular aspects of the second (KCT), a better understanding of the students' intuitions (KCS) or more explicit use of notation (SCK) could have made this example effective.

2.3.2 Mathematical Knowledge for Teaching in Upper-Division Math

As may be gathered from the example above, specifically unpacked examples of MKT are rare in the literature. In fact, we have found that many researchers develop their own sub-framework based around that of Ball et al.. For further examples, we refer the reader to Speer et al. (2010), Charalambous and Hill (2012), and McCrory et al. (2012). Not all of these are as specifically tied to undergraduate mathematics as in Section 2.3.1; however, Speer et al. (2010) outline a breakdown and preliminary analysis of teaching practices which take place at the college level and have been fruitfully analyzed at the elementary and secondary levels. Of particular note for the purposes of this project are *allocating time within lessons*, *selecting and sequencing content within lessons*, and *motivating specific content* (p. 108-9). Speer et al. observe that, at the collegiate level in particular, "the tension between rich content and limited instructional time forces teachers to make hard choices about what to include and exclude (or put off until later) and how long to spend on particular topics and activities," and that there is a rich literature on the selection of examples for college classes (p. 108). In terms of motivation, too, they remark that undergraduate instructors have the liberty of giving their motivating statement "a 'meta' character [so that] it portrays and unifies a body of content from some perspective, rather than presenting any new content" (p. 109).

Chapter 3

Mathematical Knowledge for Teaching Differential Geometry

In our study of Mathematical Knowledge for Teaching in upper-division mathematics, we examined the teaching of an introductory course in differential geometry. Differential geometry draws on many different areas of mathematics, including linear algebra, analysis, topology, and differential equations. Since there is no graduate program in differential geometry in Claremont, only advanced undergraduate mathematics students and physics students who wish to continue in the study of theoretical physics take this course. Furthermore, a teacher of differential geometry must necessarily have a grasp of the common content knowledge of these fields, and must be able to integrate and leverage the different pedagogical content knowledges from several different fields to teach differential geometry effectively. To explore how these knowledges interact in the teaching of differential geometry, we examined two texts (Do Carmo, 1976; Guillemin and Pollack, 2010) that present differential geometry at the level of an advanced undergraduate or early graduate student. We also carefully examined the slides prepared by Professor Weiqing Gu for her undergraduate "Introduction to Differential Geometry"course taught at Harvey Mudd College, and tutored and graded for this class.

3.1 Texts

Manfredo P. do Carmo's *Differential Geometry of Curves and Surfaces* is the text used in Prof. Gu's differential geometry course. It approaches differential geometry by explicating important concepts–namely, curves and surfaces–in the geometry of $\mathbb{R}^3$, a familiar space which is straightforward to visualize. Assuming that the reader has a basic understanding of the fields mentioned above, Do Carmo (1976) builds geometrically important definitions out of the fundamental ideas of derivatives, inner products, linear transformations, and homeomorphisms, and carefully illustrates how each concept contributes to the given definitions.

Guillemin & Pollack's *Differential Topology*, by contrast, begins with the much more abstract idea of a diffeomorphic embedding of $\mathbb{R}^k$ in $\mathbb{R}^n$. Abandoning consistent visualizations, Guillemin and Pollack (2010) instead focus on motivations through a generalized view of manifolds by way of Euclidean space.

Introduction to Differential Geometry is an upper-division elective mathematics course at Harvey Mudd College taught only by Professor Weiqing Gu. It is designed to introduce students to the fundamental topics studied and methods used by differential geometers. Having Multivariable Calculus, Systems of Linear Differential Equations, and Intermediate Linear Algebra as prerequisites, containing a great deal of upper-division mathematical knowledge, and being particularly applicable to physical problems, this course is intended for advanced students studying mathematics or physics. Although the course focuses primarily on the study of curves and surfaces in $\mathbb{R}^3$, it presents the material so that students can generalize the concepts to begin to understand abstract manifolds which may not be emedded in Euclidean space $\mathbb{R}^n$. However, because of the highly advanced nature of the material of this course , Professor Gu also aims to prepare the students of this class for graduate studies in mathemtics, in which concepts are presented with minimal review in class and students are expected to fill in any gaps in their knowledge or understanding through individual research. Graduate students are also expected to attend professional symposia and lectures on advanced mathematics. In addition to this, the class is bound by the normal time constraints inherent in a school year, and, as a result, the material in this class is presented with minimal scaffolding. This has a dramatic effect on the relative importances of each domain of MKT: in this pedagogical context, specialized content knowledge is minimized, while knowledge of content and students becomes especially important.

Professor Gu's lecture slides, as we would expect, closely follow Do Carmo

(1976)'s approach by focusing on definitions, theorems, and examples in $\mathbb{R}^3$. However, they make consistent references to the larger philosophy of differential geometry by discussing how each concept may be modified to apply both to surfaces in higher dimensions and to manifolds that are independent of Euclidean spaces.

3.2 Methods

When studying any aspect of teaching, and especially when studying teacher knowledge, we would ideally like to be able to observe teachers and students both in classes and during activities outside of classes that relate to the course, such as the preparation of lessons, the generation of exams, or the completion of homework. We would also like to evaluate the students' learning as a way of judging the teacher's effectiveness. However, due to external constraints, such measurements were not feasible for this project. Instead, we relied on a number of indirect measurements to examine teacher knowledge, including taking notes during tutoring sessions of student issues, conducting a number of private conversations with Prof. Gu about her own teacher knowledge, and creating our own lessons for a course in a distance-learning format.

Due to the time and space constraints of this thesis, we decided, based on our preliminary investigations and prior knowledge of differential geometry, to focus our efforts on two particular subtopics: the local properties of smooth curves and basic properties of 2-dimensional manifolds in $\mathbb{R}^3$. These topics correspond very nicely to the first two chapters of Do Carmo (1976) as well as to Lectures 2 and 5 of Prof. Gu's slides. These subtopics appear early in each of the texts and are fundamental, either by direct application or by generalization, to the further study of arbitrary manifolds. These concepts would likely be covered in any introductory course in differential geometry, and provide for rich dynamic visualizations. While they are introductory, they also are rife with subtleties and important connections which make up a significant portion of MKT.

Working both from notes from tutoring sessions and our own close readings of these texts, we generated an extensive list of questions that students had or might be expected to have with the material. This list is not exhaustive, but instead attempts to capture a representative range of topics which might cause confusion or spark interest in a student, ranging from questions about basic concepts, definitions, and mechanics to advanced insights that foreshadow future topics or relate to other fields of mathematics. We then asked Prof. Gu to respond to these questions as she would if a student were asking them in class, or to note that such a question likely would not arise in class. This process revealed both the specialized content knowledge and knowledge of content and students necessary for answering student questions in the classroom setting, which we otherwise would not have been able to measure, and helped us to evaluate more thoroughly the extent and types of student knowledge that a teacher of differential geometry assumes.

Finally, drawing on the MKT already identified, we generated lessons for each of the subtopics identified above, as well as a handful of supplemental materials, in Wolfram's Computable Document Format (CDF). This format combines elements of plain text formatting and evalutable cells of Mathematica code. Additionally, CDFs can be embedded in web pages and viewed using Wolfram's CDF Player, which is available for free. The generation of these lessons allowed us the opportunity to explore the knowledge of content and teaching that goes into creating an upper-division mathematics course, as well as that which is specific to distance-learning formats. It also helped us to further identify the SCK necessary for these contexts.

3.3 Mathematica-Based Lessons

In the creation of our own lessons on differential geometry for mass distribution, we follewed approximately the lesson plan of Prof. Gu's lectures, which correspond approximately to Do Carmo's development of the same ideas. This choice was based on practicality: the advantage of our course lies in Mathematica's ability to produce interactive graphics that help students visualize the geometry, and Guillemin and Pollack consider manifolds of arbitrary dimension, while Do Carmo begins with an extensive study of $\mathbb{R}^3$. We also chose to follow Prof. Gu's sequencing of subtopics after we determined that other sequences did not allow for appropriate introduction of terminology. In our first lesson, we introduced parametrized curves and added conditions until we could derive the Frenet Frame and introduce the Fundamental Theorem of the Local Theory of Curves. This lesson is supplemented by a mini-lesson on the differences between geometric and algebraic views of space and vectors. In our second lesson, we introduced the definition of "regular surface" given below and provided a thorough explanation of the conditions involved in that definition.

However, the main innovation in the development of these lessons was not the sequencing, but the dynamic, interactive examples provided through-out. Due to the nature of CDFs, we were restricted to `Dynamic` code; that is, code that allows interaction from the user. We relied heavily on Mathematica's `Animate[]` and `Manipulate[]` functions, but later in the process experimented with new features in Mathematica 9, including `Locator` objects and `TabView[]` windows. While there were several technical hiccups in the development of these applications, we found on the whole that the Mathematica environment is very well suited to pedagogical uses. We hesitate to label a knowledge of Mathematica functions and syntax as pedagogical content knowledge, as it is used in many situations beyond education, but we strongly recommend that teachers interested in dynamic examples, distance learning, or inverted classrooms make use of Mathematica's capabilities. All lessons described here are available at http://math.hmc.edu/seniorthesis/archives/2013/npinsky/lessons.

3.3.1 Curves

Before introducing any specific content, we provided the reader with a brief outline of the lesson to help motivate and guide the lesson. We explained that we would first define what we meant by a "curve" and then impose successive conditions on this definition until we could construct a moving frame of reference, called the Frenet Frame, on the curve. We would then use the Frenet Frame to state the Fundamental Theorem of the Local Theory of Curves.

Parametrized Curves

Our first examples, consisting of the manipulable images shown in Figures 3.1 and 3.2 along with written explanations, were used to illustrate to the reader what a differentiable parametrized curve would and would not look like, and to provide a model of how they should be visualized. Figure 3.1a shows a plane curve that is not differentiable at one point, Figure 3.1b shows a plane curve that is differentiable everywhere but contains a self-intersection, and Figure 3.2 shows a curve in $\mathbb{R}^3$ that is differentiable everywhere and does not contain self-intersections (in fact, the curve in Figure 3.2 is parametrized by arc length, and we return to this example later in the lesson). Each of these examples contain sliders that move a point along a trace, helping the reader to think of the curve as the function by visually identifying a point in the domain (the slider) with a point in the range (marked by a small disk). Taken together, these examples illustrate a range of ways in which a function may fulfill or fail to fulfill the definition

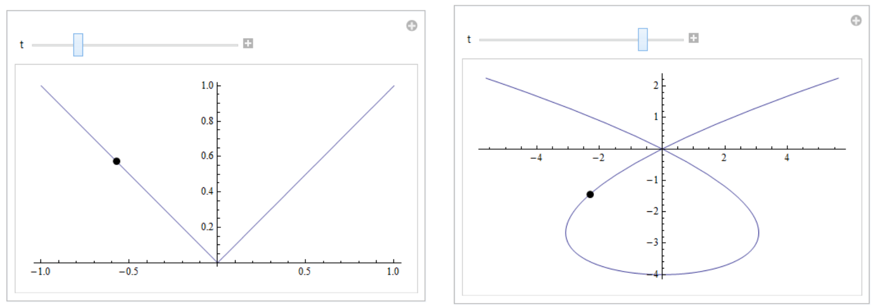

a. A parametrized curve that is not differentiable at one point.

b. A parametrized curve that is differentiable everywhere.

Figure 3.1 Interactive illustrations of parametrized curves in the plane, which highlight differentiability and allow the reader to traverse each curve's trace.

of a differentiable parametrized curve.

To emphasize the fact that we define a curve to be a function and not the trace of a function, we provide the manipulable shown in Figure 3.3, which allows the reader to watch as two different curves α and β trace out the unit circle at different rates. Both curves are functions of the same parameter t, which allowed us to attach both to the same slider. Thus, the reader can see that the same point in the domain corresponds to different points in the range under the two different curves, further helping distinguish the functions from their traces.

To introduce the concept of a regular curve, i.e., a curve with a nonvanishing first derivative, we provide two examples of non-regular curves in Figure 3.4. Figure 3.4a gives a curve which is differentiable everywhere–it is given by $t \mapsto (t^3, t^2)$–but has a cusp and a singular point at $t = 0$. The plot shows both a point, representing the function itself, and a vector based at that point, representing the first derivative of the function. The reader can watch as the vector shrinks to 0 as the point approaches the origin from either direction. However, so that the reader does not mistakenly believe that singular points occur only at cusps or corners, we include an example of a curve which has a straight line as its trace but has a singular point (Figure 3.4b). The limitation of this example, though, is that it does not help illustrate why singular points may prevent parametrization by arc length. This concept could be more effectively illustrated by a curve that doubles back on itself; this was not implemented due to time constraints.

Regularity is important in the context of differential geometry because it allows us to reparametrize by arc length. To help explain this process,

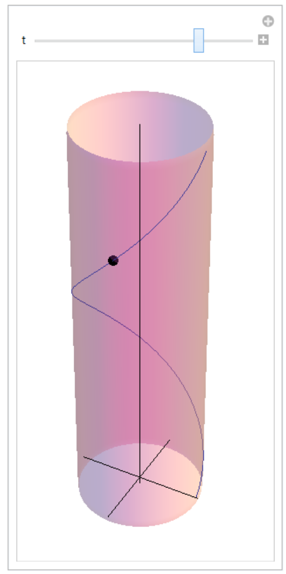

Figure 3.2 Interactive illustration of a parametrized curve in $\mathbb{R}^3$, with a semi-transparent cylinder to help orient the reader.

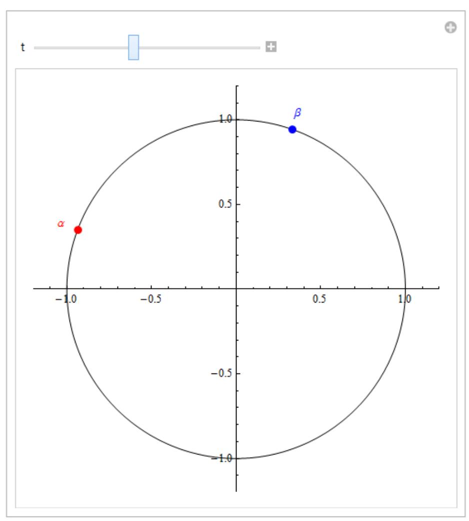

Figure 3.3 An interactive illustration of the fact that different curves may have the same trace.

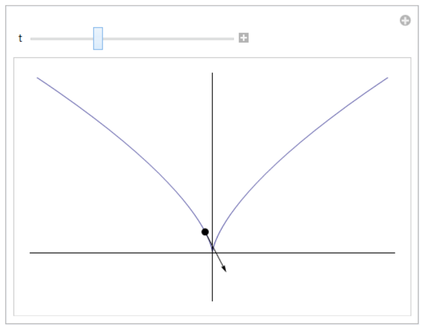

a. A curve that is differentiable, despite its trace having a cusp. At the cusp, the curve has a vanishing derivative, and is therefore not regular.

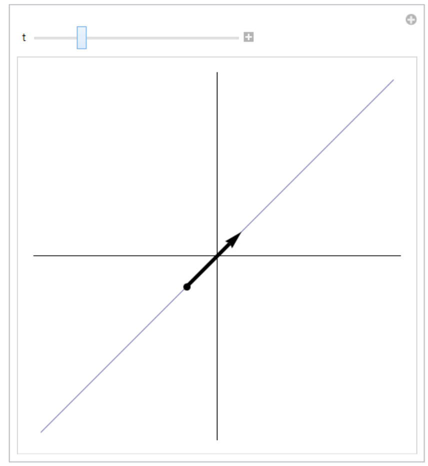

b. A pathological case in which a line is parametrized so that it is not regular. The derivative vanishes at the origin.

Figure 3.4 Interactive examples of curves that are *not* regular, helping the reader to visualize such curves

Figure 3.5 illustrates the computation of the arc length of a curve. Using the slider, the reader can trace out the curve as in previous example, but this example highlights the portion of the curve that has been traced using a thick red line. Above the graphic, a dynamically updated display shows the arc length, computed numerically, of the red line. As the curve has nonconstant speed, the reader can see that the arc length depends only on the trace of the curve and not on the parameter.

The Frenet Formulae

Parametrization by arc length, toward which the examples in the previous section built, plays a fundamental role in the differential geometry of curves, as it ensures the orthogonality of the first and second derivatives of a curve. This is illustrated by the manipulable shown in Figure 3.6a, which shows the two curves from Figure 3.3 along with their first and second derivatives. In this instance, however, the curves have different sliders to emphasize that they are parametrized differently and to allow the reader to compare the first and second derivatives at the same point in the trace. The orthogonality exhibited in Figure 3.6a is then used construct the Frenet Frame, which is illustrated in Figure 3.6b. In this example, we include a

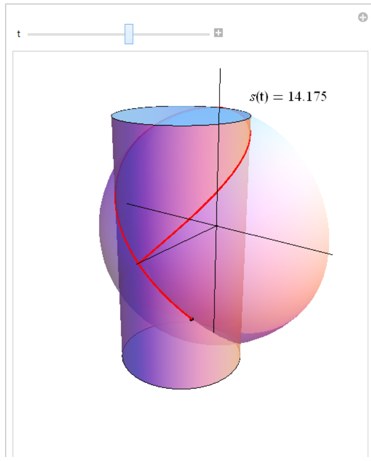

Figure 3.5 A manipulable that allows the reader to watch how arc length is computed as the trace of a curve is traversed.

bounding box to help the reader orient themself, and draw the tangent, normal, and binormal vectors in different colors to help the reader track them. It is important to note here that the normal vector (shown in blue) is not the second derivative of the function, but rather a unit vector in that direction. Since the second derivative of the red curve in Figure 3.6a is already a unit vector, we were concerned that readers might make this mistake, and took care to emphasize the distinction in writing and to illustrate it with Figure 3.7a.

Curvature and torsion are central to the establishment of the Fundamental Theorem of the Local Theory of Curves. To help the reader envision them, we included the two examples shown in Figure 3.7. In order to illustrate curvature, we selected the plane curve shown in Figure 3.7a, which displays the trace of the curve, the tangent, normal, and second derivative vectors, and the osculating circle. We chose to use a plane curve to help the reader focus on the curvature and osculating circle without being distracted or disoriented by three-dimensional behavior. As the curve spirals in to the origin, the reader can focus on how the osculating circle approximates the curve and how its radius is related to the length of the second derivative vector (shown in red). Plane curves also allow us to draw the Frenet Frame without parametrization by arc length, as we may simple apply a rotation to the tangent vector to obtain the normal vector. Because reparametrization by arc length is computationally intensive, this significantly broadened our choice of curves for this example.

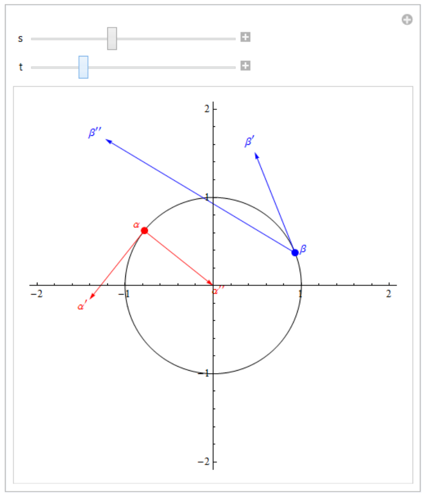

a. An example illustrating the benefits of parametrization by arc length.

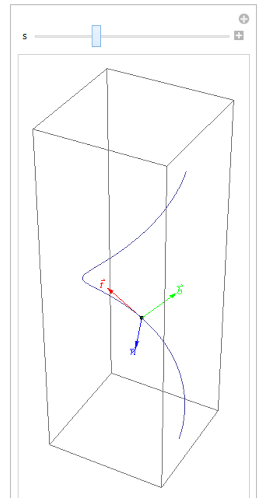

b. An example illustrating the Frenet Frame on a helix.

Figure 3.6 Two examples illustrating how parametrization by arc length allows us to establish an orthonormal basis for a neighborhood of a point on the curve.

On the other hand, we were forced to use a space curve to illustrate torsion because plane curves have zero torsion, and were therefore forced to search for a curve that was either already parametrized by arc length, or easily reparametrizable, so that we could draw the Frenet Frame on the curve. The result is shown in Figure 3.7b, which allows the reader to traverse the curve while simultaneously observing the behavior of the binormal vector's derivative and of the torsion function. Furthermore, the reader is able to rotate the 3D plot to more carefully examine how the curve itself "pulls away" from the osculating plane (also shown), and how the binormal vector moves as this occurs.

The Fundamental Theorem of the Local Theory of Curves

Finally, to illustrate the Fundamental Theorem of the Local Theory of Curves, which is the culmination of our lesson on curves, Figure 3.8 shows two curves which differ by a rigid motion. As with previous curves, we include one slider to allow the reader to traverse the curves with the Frenet Frame. To emphasize that these curves really do differ only by a rigid motion, the same parameter is used for both. A second slider is also provided to allow the reader to apply a rigid motion to one of the curves, bringing it to align

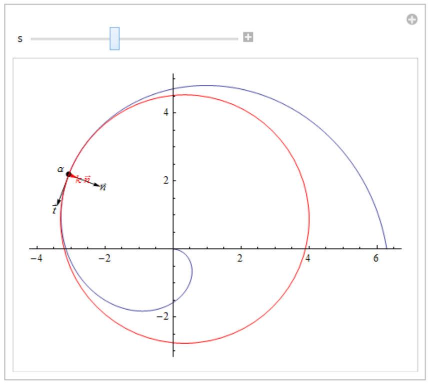

a. The curvature of a curve and its relationship to the radius of the osculating circle.

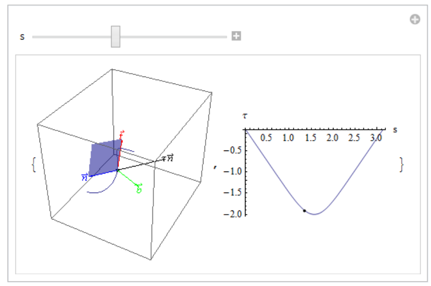

b. Torsion and its relationship to the 3D behavior of a curve.

Figure 3.7 Interactive examples illustrating curvature and torsion in curves.

with the other. We had hoped to include plots of the curvature and torsion for these curves to further emphasize that they are the same, but this proved too computationally intensive.

Views of Space

As a supplement to our lesson on regular curves, we created a mini-lesson on how differential geometers view vectors. The mini-lesson presented both the algebraic perspective, in which vectors are based at a pre-established origin and expressed in terms of coordinates, and the geometric perspective, in which vectors have no inherent base point, but one may be chosen to establish a coordinate system. We also explained how these two perspectives are reconciled by differential geometers by considering the parallel transport of a vector to be based at the origin. Figure 3.9 shows the interactive examples in this mini-lesson. Figure 3.9a allows the reader to uniquely determine a vector in $\mathbb{R}^3$ by manipulating the coordinates, while Figure 3.9b lets the reader drag the base point and vectors around the plane, and draws the resulting coordinate system. Figure 3.9c reconciles these perspectives by visually redrawing the reader's vector so that it is based at the origin.

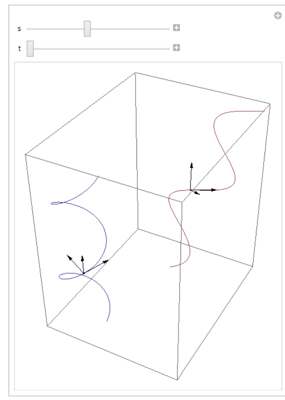

Figure 3.8 Interactive example allowing the reader to move one curve "on top of" the other via a rigid motion. The two curves have the same curvature and torsion, and therefore align after the rigid motion.

3.3.2 Surfaces

Our lesson on regular surfaces began by motivating the study of such surfaces via a discussion of state spaces for physical systems. We then provided Do Carmo (1976)'s definition for a "regular surface" and a detailed explanation of each of the three conditions involved.

Since we expect most students entering differential differential geometry not to be familiar with homeomorphisms, we provide several visual illustrations of this concept, shown in Figure 3.10. These examples, presented in tabs in a single window, attempt to provide the student with a mental storehouse of sets in $\mathbb{R}^3$ that either are homeomorphic to the plane or fail to be homeomorphic in very specific ways. For this reason, we include sets that are homeomorphic, but not diffeomorphic to the plane, sets for which a function from the plane could not be continuous, and sets for which a function into the plane could not be continuous.

We also expect that most students will be unfamiliar with the differential of a map, and therefore explain it and its relationship to the definition of a regular surface extensively. First, we provide the example shown in Figure 3.11a, which allows the reader to move a vector around the domain of a parametrization and to observe how the differential maps this vector to $\mathbb{R}^3$. In this example, the use of `Locator` objects is especially valuable, as it permits the manipulation of a point in $\mathbb{R}^2$ without a separate slider for each coordinate. Figure 3.11b then provides a pathological example of a parametrization which is not regular along a line, despite its image being

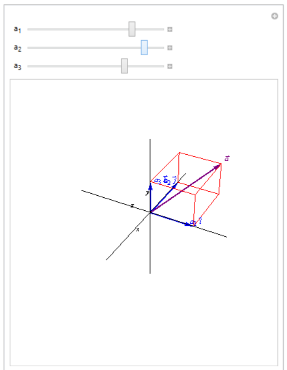

a. Interactive example of the algebraic perspective of vectors, allowing the reader to manipulate the three components individually.

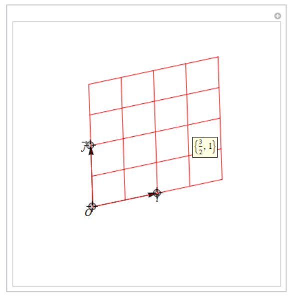

b. Interactive example of the geometric perspective of vectors, allowing the reader to move the origin and basis vectors independently.

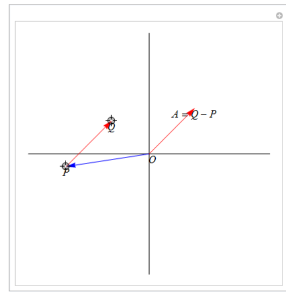

c. Interactive example of how we reconcile the algebraic and geometric perspectives of vectors, illustrating the parallel transport of a vector to the origin.

Figure 3.9 Interactive examples of the various views of vectors in Euclidean space.

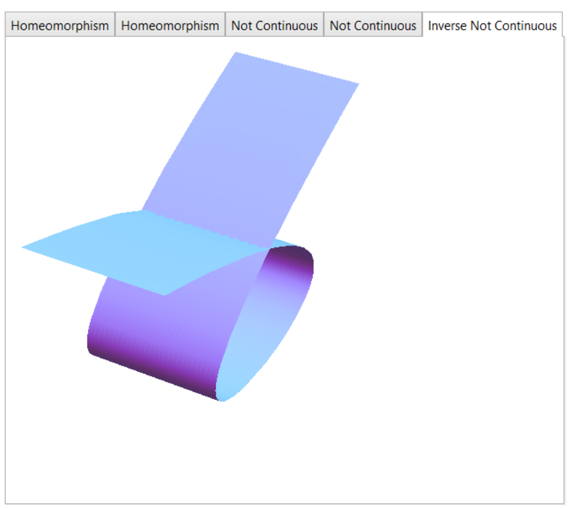

Figure 3.10 A set of examples illustrating the concept of homeomorphic sets, and the various ways in which a set can fail to be homeomorphic to a plane.

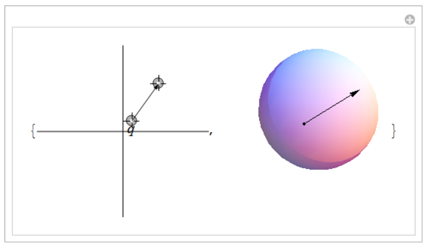

a. Allows the reader to manipulate a vector in the domain of a function and observe where it is mapped to.

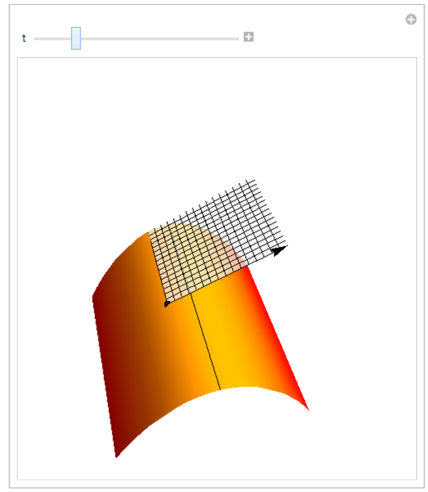

b. An illustration of a pathological case in which a parametrization is not regular at all points, despite the surface having a well-defined tangent plane at those points.

Figure 3.11 Interactive examples illustrating the behavior of the differential of a parametrization.

a regular surface and having a well defined tangent plane at all points of this line. By showing the tangent plane shrinking to a single dimension as we approach this line, we emphasize to the reader how the differential of a map lets us discuss symbolically the tangent plane.

We would like to reemphasize that the techniques used in the generation of our lessons are not knowledge used exclusively for teaching, but rather techniques from other fields that are particularly relevant for our use.

3.4 Results

In our exploration, we relied heavily on the specific tasks identified by Ball et al. (2008) as aligned with each subdomain of MKT. Unsurprisingly, we found that many of these tasks are present in the teaching of differential geometry. For example, an instructor must sequence topics for instruction, select pedagogically useful examples, and respond to student questions as much as or more than a grade school teacher. Nevertheless, there are several differences in the types of SCK, KCS, and KCT that an instructor of differential geometry must have and use in the act of teaching. These differences arise from the vastly different structure of a college class as compared to a class in grade school, and the differences in student-instructor relationships and interactions. Most notably, differential geometry is an elective course, meaning that students have a choice in whether they enroll in it, and whether they drop it later. This also means that the instructor can set prerequisites, guaranteeing that the students have certain knowledge. Additionally, the student-instructor relationship in a college course puts most of the responsibility for learning on the student. Both of these aspects, we note, apply as much to a distance-learning course as to a traditional in-classroom experience. Finally, the concepts taught in differential geometry are, in general, more theoretical than methodological. As detailed below, these factors greatly affect the mathematical knowledge for teaching differential geometry.

3.4.1 Specialized Content Knowledge

In the context of a class designed to encourage independent investigation and processing of concepts, concepts are, in general, presented simply, as if to a mature mathematician. As a result, much of the content knowledge that goes into creating and teaching this course is purely common content knowledge, such as the knowledge of the important definitions and theorem statements in differential geometry. However, as this is a course for undergraduate students, specialized content knowledge cannot disappear entirely. From the two topics studied in our investigation, we have concluded that specialized content knowledge in differential geometry is, in general, a decompressed form of common content knowledge. It takes three general forms: understanding the need for each hypothesis in a definition or theorem, knowing how concepts in specific contexts generalize or appear in other contexts, and knowledge for the development of pedagogically strategic examples.

As in most mathematical fields, differential geometry relies very heavily on precise definitions. In the two topics of our investigation, curves and regular surfaces, the selection of the "right" definition is essential to productive study of these mathematical objects. Although Guillemin and Pollack do not deal with curves in great depth, Do Carmo devotes the entire first chapter to them. In his presentation, Do Carmo defines curves as functions, and defines a nested sequence of types of curves:

1. A *parametrized differentiable curve* is a differentiable map $\alpha : I \to \mathbb{R}^3$ of an open interval $I = (a, b)$ of the real line $\mathbb{R}$ into $\mathbb{R}^3$.

2. We say that $s \in I$ is a *singular point of order 0* of a parametrized differentiable curve $\alpha : I \to \mathbb{R}^3$ if $\alpha'(s) \neq 0$. We say that α is *regular* if it has no singular points of order 0.

3. A regular parametrized differentiable curve $\alpha : I \to \mathbb{R}^3$ is said to be *parametrized by arc length* if $\|\alpha'(t)\| = 1$ for all $t \in I$.

4. We say that $s \in I$ is a *singular point of order 1* if $\alpha''(s) = 0$. If α has no singular points of order 1, we say that it is *regular curve of order 1*.

Properties of these curves are used to define the Frenet Frame, a moving reference frame defined from the first and second derivatives of a curve, and to state the Fundamental Theorem of the Local Theory of Curves. While these definitions and theorems might be common knowledge to any differential geometer, a teacher of differential geometry must know explicitly, for example, that the definition of a curve as a map excludes the possibility of identifying a curve with its trace, that regularity (of order 0) allows the use of the Inverse Function Theorem to reparametrize the curve by arc length, and that regularity of order 1 ensures that the Frenet Frame is well-defined. A teacher must also know, however, that isolated singular points or isolated points of non-differentiability do not prevent us from studying these curves, as we can use piecewise parametrizations for all other points.

A further example of definition-related specialized content knowledge arises from the study of regular surfaces in $\mathbb{R}^3$. Do Carmo defines a regular surface as

> a subset $S \subset \mathbb{R}^3$ [such that] for each $p \in S$, there exists a neighborhood $V \in \mathbb{R}^3$ and a map $\mathbf{x} : U \to V \cap S$ of an open set $U \subset \mathbb{R}^2$ onto $V \cap S \subset \mathbb{R}^3$ such that
>
> 1. $\mathbf{x}$ is differentiable,
> 2. $\mathbf{x}$ is a homeomorphism, and
> 3. $\mathbf{x}$ is regular [i.e., the differential $d\mathbf{x}_q$ is injective for each $q \in U$].

Again, this definition might be common knowledge, at least in a very compressed form, to a professional mathematician, but the purposes of each part of the definition are specifically specialized content knowledge. In this case, the teacher must know that this definition specifies a surface as a set covered by the images of maps, rather than as a single parametric function as in the definition of a curve, because even simple surfaces, such as the sphere, are not diffeomorphic, or even homeomorphic, to the plane. Specialized content knowledge also includes the knowledge that condition 1 in the definition allows differential calculus and condition 2 allows the use of analysis and topology in the study of the surface, while condition 3 allows the use of linear algebra for the study of the tangent plane of the surface.

The second type of specialized content knowledge, knowledge of how concepts generalize or appear in other contexts, also arises in the study of these two topics. In particular, both curves and regular surfaces can be viewed as types of manifolds, an understanding of which is one goal of this course. However, a teacher must know that the definitions of curves given above do not generalize to manifolds, but rather the definition of a manifold may be narrowed so as to allow a study of curves. Instead, we approach manifolds as we did regular surfaces, by covering them with the images of parametrizations. Should a student ask why, then, we define curves as we do, a teacher must be able to explain that this approach is more helpful, as it affords the construction of the Frenet Frame, and that parametrized curves are frequently used in the study of abstract manifolds.

The third type of specialized content knowledge, knowledge for the development of pedagogically strategic examples, is the most related to the other two domains of MKT discussed here. In this class, examples are used for two primary purposes: the illustration of a particular definition or theorem, and as a counterexample to dispel common student misconceptions. In the service of the first purpose, the understanding of the importance of each part of a definition or theorem comes to bear, while in the service of the second, knowledge of content and students and knowledge of content and teaching are particularly relevant. However, we can identify further knowledge that is necessary for the development of pedagogically useful examples. When teaching about curves, for example, a teacher might want to explicitly reparametrize a curve by arc length. Although any parametrized curve with nonvanishing first derivative can be reparametrized by arc length, the details of these computations are frequently extremely messy, and the arc length function may not have an inverse with a closed form. This specific example was particularly relevant in the creation of our distance-learning lesson on curves, in which we frequently needed curves parametrized by arc length in order to illustrate the orthogonality of the first and second derivatives of such curves. We quickly built up a storehouse of such curves for reuse in the lesson.

3.4.2 Knowledge of Content and Students

In their discussion of Knowledge of Content and Students, Ball et al. suggest that this domain of teacher knowledge includes the processing of student ideas, familiarity with common errors, deciding which of several errors students are most likely to make, predicting what students will find interesting and motivating, and anticipating the level of difficulty students will have with a given task. Of these, we discard the processing of student ideas and the prediction of what students will find interesting and motivating. We do not discard the former because it is not important in such a class; to the contrary, this is one of the most imporant things a teacher does, and what distinguishes a teacher from a textbook. We discard it simply because, as noted above, we had no opportunities to observe in-classroom interactions during the course of this project. We discard the prediction of what students will find interesting and motivating because, whether we are considering a traditional or distance-learning classroom, the course is elective and high-level, and we assume that students in such a course are already interested in the material and motivated to learn it. We will address each of the remaining tasks in turn.

As in all classes, familiarity with common errors or misconceptions is vital to effective teaching. In a college or online course setting, however, a teacher does not have the same access to student work as an elementary or secondary teacher does. Instead of grading regular homeworks themself, these teachers frequently have graders or crowd-source the grading of homework to the members of the class, and exams in these settings are far less frequent than in elementary or secondary school–instead of one after every chapter in a textbook, for example, a college or online class instructor will likely schedule only a midterm exam and a final exam. A teacher's knowledge of common errors, therefore, is important primarily during lectures or class discussions, when students raise questions. In order to use class time most effectively, a teacher will likely elect to address such misconceptions in class. In differential geometry, one common misconception arises from the definition of a regular surface given above. Although the definition specifies that "for each $p \in S$, there exists a neighborhood $V \in \mathbb{R}^3$ and a [parametrization] $\mathbf{x} : U \to V \cap S$" that satisfies three conditions, students can misunderstand some or all of the definition and believe that each parametrization must cover all of S, that every map from U to $V \cap S$ must satisfy the three conditions, that neighborhoods $V \cap S$ and $W \cap S$ belonging to different parametrizations must be disjoint, or any number of other possible errors. As students are likely to ask questions about these types of misunderstandings in class, it is worth the teacher's time and effort, even in a class designed around independent learning, to address these issues. For example, Prof. Gu's class uses an example from Do Carmo in which the sphere is covered with the graphs of functions of the form $x_3(x_1, x_2) = \pm\sqrt{x_1^2 + x_2^2}$, where x_1, x_2, x_3 are the x, y, and z coordinates in some order. Prof. Gu also makes a special point of the difference between the definitions of a curve as a map and a surface as a set covered by maps in order to help students make the mental transition. In this process, the teacher not only must have familiarity with common student errors, but must also decide which ones students are likely to make in order to address these in class.

More important, though, in the context of a class that emphasizes independent student learning is the teacher's ability to judge the difficulty level of different tasks that students are expected to accomplish. This is clearly an issue in teaching at any level, as all teachers assign exercises for students to practice and assessments to gauge their learning. In this course, however, students are also expected to pursue gaps in the material presented in the classroom, and so a teacher must take great care when gauging where to leave these gaps, and where to fill them in with explanations. In the lessons examined, this is most apparent in the explanation of condition 3 in the definition of a regular surface. Condition 3, which is stated as "$\mathbf{x}$ is regular," requires that the differential $d\mathbf{x}$ of a parametrization $\mathbf{x}$ be injective at all points in its domain. Recognizing that the definition of the differential does not make immediately clear what injectivity would look like, Do Carmo (1976) and Prof. Gu give the differential a more familiar form by showing that the matrix of the differential with respect to the canonical bases is the Jacobian. This then allows the observation that the differential is one-to-one if and only if the columns $\partial\mathbf{x}/\partial u$ and $\partial\mathbf{x}/\partial v$ are linearly independent, or equivalently, if any of the Jacobian determinants

$$\frac{\partial(x,y)}{\partial(u,v)} = \begin{vmatrix} \frac{\partial x}{\partial u} & \frac{\partial x}{\partial v} \\ \frac{\partial y}{\partial u} & \frac{\partial z}{\partial v} \end{vmatrix}, \qquad \frac{\partial(x,z)}{\partial(u,v)}, \qquad \frac{\partial(y,z)}{\partial(u,v)}$$

are nonvanishing. While these pieces of knowledge may be considered specialized content knowledge, the knowledge that students would not likely think to compute the matrix of the differential (in any basis) is very clearly knowledge of content and students. Similarly, the knowledge of the relationships between these equivalent statements may be considered specialized content knowledge (or perhaps common content knowledge), but the understanding that students can deduce these relationships for themselves solely from what is given allows the teacher to avoid further explanation, and therefore is knowledge of content and students.

3.4.3 Knowledge of Content and Teaching

In the given pedagogical context, two aspects of Knowledge of Content and Teaching suggested by Ball et al. are particularly notable. The first, knowledge for sequencing topics for instruction, does not at first glance seem as though it would be important: differential geometry has a very natural progression from one-dimensional curves to two-dimensional surfaces to n-dimensional manifolds. However, there is considerable flexibility in the order in which these concepts are presented, as well as in the subtopics contained in each. The second, knowledge for the evaluation of multiple representations, is less interesting than in elementary or secondary classes. Given the mathematical maturity of students capable of taking differential geometry and the emphasis on independent learning in the class, it is rarely advantageous to present only a select number of representations, as it is frequently a worthwhile exercise for students to connect these representations themselves. Nevertheless, there are a few scenarios in which a teacher of differential geometry does need to weigh the merits of different representations, which we discuss below.

The ultimate goal of this introduction to differential geometry is for students to understand the fundamental concepts of and the methods used to study abstract manifolds. There is, therefore, a sensible progression from curves (one-dimensional manifolds which are also used to study higher dimensional manifolds) to regular surfaces (two-dimensional manifolds in 3-dimensional Euclidean space) to abstract manifolds of arbitrary dimension that may or may not be embeddable in Euclidean space. Do Carmo focuses on the first two steps of this generalization, and Prof. Gu generally follows this outline, making regular mention of abstract manifolds, especially Lie groups like $SO(3)$, and alluding to how concepts from the study of regular surfaces will generalize to such manifolds. Guillemin and Pollack, however, approach differential geometry from the opposite angle, beginning with the definition of a diffeomorphism and then moving to an examination of manifolds in $\mathbb{R}^n$. In either of these approaches, a teacher must recognize that some concepts, such as diffeomorphisms and differentials, are vital and must appear early in the presentation of the material. Both approaches are viable, and a decision between them will necessarily involve knowledge of content and students and a clear goal for the course. Making this choice, though, will necessarily affect the content that may be covered in the course. For example, although Guillemin and Pollack's approach gives an immediate broad understanding of manifolds that may be embedded in $\mathbb{R}^n$ and eventually allows them to discuss transversality of manifolds, it may prevent students from holding a view of a manifold without reference to an ambient space, as both Gu and Do Carmo do.

Decisions about the sequencing of topics can also take place within lessons, as we found in the creation of our lesson on regular curves. When we began work on this lesson, we saw an alternative to the arc used by Do Carmo and Prof. Gu. Instead of presenting the most basic definitions of curves and successively narrowing the definition to obtain useful properties, such as the orthogonality of the first and second derivatives or the relationship between the length of the second derivative and the radius of the osculating circle, eventually culminating in a statement of the Fundamental Theorem of the Local Theory of Curves, we considered beginning our lesson with this theorem and then determining what would be necessary to prove it. Ultimately, we found a compromise: we gave a general overview of the goals of the lesson (defining curves, the Frenet Frame, and the Fundamental Theorem) and referenced these goals as we followed the lesson plan given by Do Carmo and Prof. Gu. Through this choice, we hoped to effectively motivate the definitions we presented while still providing adequate scaffolding for student learning.

The second facet of Knowledge of Content and Teaching, knowledge for the evaluation of multiple representations, arises in upper-division mathematics when there are multiple equivalent definitions for a concept, or multiple equivalent statements of a condition. One such example is the regularity condition (condition 3) in the definition of a regular surface. Although it is initially given as "the differential $d\mathbf{x}$ is injective at all points in the domain of $\mathbf{x}$," both Prof. Gu and Do Carmo derive the matrix of $d\mathbf{x}$ with respect to the canonical bases and point out several conditions equivalent to injectivity, as described above, but do not explain why these conditions are equivalent. It is left to the student as an exercise to prove such equivalence, or at least justify it to themselves.

Another example of this type of evaluation occurs when defining the differential of a map just prior to the discussion of the regularity condition. In our sources, we found two equivalent definitions for the differential. Do Carmo and Prof. Gu, who deal specifically with subsets of $\mathbb{R}^3$ that are locally diffeomorphic to $\mathbb{R}^2$, define the differential in terms of the derivatives of corresponding curves in the plane and in the surface, while Guillemin and Pollack, who deal more generally with subsets of $\mathbb{R}^n$ that are diffeomorphic to $\mathbb{R}^k$, define the differential simply as a directional derivative. As with sequencing, the ultimate valuation of these decisions depends on the ultimate goal of the teacher. Prof. Gu aimed to intoduce her students to manifolds like $SO(3)$ that do not naturally live in Euclidean space, and therefore needed a definition that could generalize to such manifolds. Guillemin and Pollack, on the other hand aimed to focus primarily on manifolds that could be embedded in $\mathbb{R}^n$, and could therefore use a definition that relied on $\mathbb{R}^n$. In our lesson, we presented both definitions and noted their equivalence, as we felt that Guillemin and Pollack's definition is easier to visualize for a first-time student, while Do Carmo's definition more readily generalizes to abstract manifolds. As Prof. Gu and Do Carmo did with the regularity condition, we left the details of this equivalence to the student so that they might solidify their understanding on their own.

3.4.4 Knowledge of Content and Curriculum

Due to the elective nature of the class noted above, an instructor in an upper-division mathematics course has greater control over incoming student knowledge. In setting prerequisites, a college professor can specify the content knowledge that students enter the class with, as well as predetermine the methods they know and the philosophical approach that has guided their learning. In a private conversation with Prof. Gu in which we discussed the list of questions we might expect a student to have, she made frequent references to the prerequisite material for the course in providing answers to questions or explanations of processes and principles. In fact, the bulk of the questions discarded as unlikely to arise in a normal class were explained to be part of the prerequisite material, understood tacitly by both the students and professor. Of course, this is not to say that students will enter a course with uniform comprehension of the prerequisite material, or that they will remember perfectly or recognize immediately the aspects of the prerequisite material most pertinent to differential geometry (or any other course). However, the instructor, having set the prerequisite knowledge, can draw on those principles that the students have already learned or refer the students to theorems that may have slipped from their memory. A vital piece of a college mathematics instructor's knowledge, therefore, is a knowledge of the specific curricula of prerequisite courses.

Chapter 4

Conclusion

In our analysis in the preceding chapter, we identified a number of items of mathematical knowledge for teaching, classified by subdomain, that are important for the teaching of differential geometry. In general, these items are not qualitatively different from what we would expect based on extrapolation from the work of Ball et al.. Quantitatively, however, they differed greatly, both in absolute number and relative importance. Specifically, specialized content knowledge seemed to be far less important that knowledge of content and students, as all examples of the former dealt with particular topics within the class, while the latter focused on overall approach to structuring and teaching the class. These differences seem to be due primarily to the secondary pedagogical goal of Prof. Gu's course (preparation for graduate school studies) and the mathematical maturity of the students. By focusing on preparation for graduate studies, a teacher necessarily will cut out sections of instruction that would otherwise be included, so as to force the students to research and make connections on their own. This is related to a recognition of students' ability to accomplish these tasks, especially making sense of new concepts and notations and relating them to previously known ideas. Indeed, we believe that, were it reasonable to teach differential geometry to elementary-level students, the mathematical knowledge for teaching involved in this course would align both qualitatively and quantitatively with the mathematical knowledge for teaching a subject like multiplication or simple algebra.

To understand more clearly the philosophical underpinnings of this pedagogical goal, we recall Ball et al.'s explanation for the need for mathematical knowledge for teaching:

> Teaching involves the use of decompressed mathematical knowledge that might be taught directly to students as they

> develop understanding. However, with students the goal is to develop fluency with compressed mathematical knowledge. In the end, learners should be able to use sophisticated mathematical ideas and procedures. Teachers, however, must hold unpacked mathematical knowledge because teaching involves making features of particular content visible to and learnable by students. . . . Teachers [also] must be able to talk explicitly about how mathematical language is used; how to choose, make, and use mathematical representations effectively; and how to explain and justify one's mathematical ideas. All of these are examples of ways in which teachers work with mathematics in its decompressed or unpacked form (Ball et al., 2008: p. 400).

The "decompressed mathematical knowledge" that Ball et al. discuss is the same knowledge that they classify as mathematical knowledge for teaching, and arrange into the subdomains presented in Chapter 2. However, the tasks they present are tasks that we would expect any professional mathematician to be able to perform; in fact, these are tasks that mathematicians are required to engage in as a part of their daily work. It is preposterous to imagine, for example, a conference presentation in which the presenters do not justify their mathematical ideas, use mathematical representations effectively, or make clear the language they are using.

We suggest, therefore, that differences in MKT found in our differential geometry course when compared to that of Ball et al. are more revealing of the expectations placed on students than of aspects of the material that alter how it is taught. In grade school, teachers have very little idea as to what careers their students might enter. They therefore must teach the most practical mathematics, which they may apply in whatever profession they choose or which is generally accepted for basic mathematical literacy. In college and beyond, however, students frequently have a profession or field in mind, and the mathematics may be catered to their needs. If that field lies in mathematics research, then they need a decompressed mathematical knowledge in order to engage with colleagues. This in turn can change a teacher's role from instructor to veteran guide.

These ideas about pedagogical goals are not new to the field of education. In a discussion of the failure of U.S. public schools, Labaree (1997) asserts that in any educational context, a teacher must make decisions about what content to cover and what aspects of that content to emphasize. Since the process of setting these goals "is resolved through a process of making choices. . . [based on] values and interests," it is fundamentally a political issue, not a pedagogical one (40). Labaree identifies three distinct philosophical and political goals of public education in the United States: democratic equality, social efficiency, and social mobility. Democratic equality refers to the effective training of citizens to participate in the democratic process, social efficiency to the training of skilled but differentiated workers to power the economy, and social mobility to the development of skills for personal advancement in society. Although Labaree's analysis referred to large-scale educational policy, these political orientations affect classroom interactions by dictating a teacher's curriculum.

Labaree's political goals, however, do not capture the philosophical-political goals of either college or distance learning courses, as these may not be considered "public education." Labaree's argument refers specifically to public education, in which the government pays for the education of the country's youth. Since the government runs the education system, it demands goals that will benefit society as a whole and requires every child to take part in education. This is, however, not the case in either higher education or independent distance education. In both cases, education is pursued voluntarily and paid for by the students. Furthermore, students have a great deal more discretion as to which upper-division elective classes they take, and can choose to drop a course mid-semester. Teachers, too, have a great deal more autonomy in the curricula and instruction of these courses. Instead of being bound to state or federal regulations or catering to the core educational goals of their school or university, these teachers can design a course precisely to their specifications, teaching exactly what they find interesting or useful in exactly the way they choose. We suggest that Labaree's ideas of political goals for teaching might be applied to help gauge "good teaching" in different pedagogical contexts.

In this light, the mathematical knowledge for teaching differential geometry, which deemphasizes specialized content knowledge and emphasizes knowledge of content and students, takes the form it does specifically because of the philosophical-political values of the course. More generally, we hypothesize that the mathematical knowledge for teaching any course across the educational spectrum, regardless of the content knowledge, is shaped in large strokes by the goals of the teacher, department, school system, state, and nation in which it is located, and through these goals, teaching is influenced by the values of the ambient society.

Bibliography

Arnold, Douglas N., and Jonathan Rogness. 2008. Möbius transformations revealed. *Notices of the AMS* 55(10):1226–1231.

Ball, Deborah Loewenberg, and Hyman Bass. 2003. Toward a practice-based theory of mathematical knowledge for teaching. In *Proceedings Of The 2002 Annual Meeting of the Canadian Mathematics Education Study Group*, eds. B. Davis and E. Simmt, 3–14. Edmonton, Alberta, Canada: Canadian Mathematics Education Study Group (Groupe Canadien d'étude en didactique des mathématiques).

Ball, Deborah Loewenberg, Mark Hoover Thames, and Geoffrey Phelps. 2008. Content knowledge for teaching: What makes it special? *Journal Of Teacher Education* 59(5):389–407.

Beswick, Kim. 2012. Teachers' beliefs about school mathematics and mathematicians' mathematics and their relationship to practice. *Educational Studies in Mathematics* 79:127–147.

Charalambous, Charalambos Y., and Heather C. Hill. 2012. Teacher knowledge, curriculum materials, and quality of instruction: Unpacking a complex relationship. *Journal of Curriculum Studies* 44(4):443–466.

Choi, Hee Jun, and Scott D. Johnson. 2005. The effect of context-based video instruction on learning and motivation in online courses. *The American Journal Of Distance Education* 19(4):215–227.

Do Carmo, Manfredo Perdigao. 1976. *Differential Geometry Of Curves And Surfaces*, vol. 2. Englewood Cliffs: Prentice Hall.

Guillemin, Victor, and Alan Pollack. 2010. *Differential Topology*, vol. 370. American Mathematical Society.

Kuntze, Sebastian. 2012. Pedagogical content beliefs: global, content domain-related and situation-specific components. *Educational Studies in Mathematics* 79:273–292. URL http://dx.doi.org/10.1007/s10649-011-9347-9. 10.1007/s10649-011-9347-9.

Labaree, David F. 1997. Public goods, private goods: The american struggle over educational goals. *American Educational Research Journal* 34(1):39–81.

Larkin, Douglas. 2012. Misconceptions about "misconceptions": Preservice secondary science teachers' views on the value and role of student ideas. *Science Education* 96(5):927–959.

Lewin, Tamar. 2012. College of future could be come one, come all. http://www.nytimes.com/2012/11/20/education/colleges-turn-to-crowd-sourcing-courses.html?pagewanted=all.

McCrory, Raven, Robert Floden, Joan Ferrini-Mundy, Mark D. Reckase, and Sharon L. Senk. 2012. Knowledge of algebra for teaching: A framework of knowledge and practices. *Journal for Research in Mathematics Education* 43(5):585–615.

Niess, M.L. 2005. Preparing teachers to teach science and mathematics with technology: Developing a technology pedagogical content knowledge. *Teaching and Teacher Education* 21(5):509 – 523. doi:10.1016/j.tate.2005.03.006. URL http://www.sciencedirect.com/science/article/pii/S0742051X05000387.

Ryve, Andreas, Per Nilsson, and John Mason. 2012. Establishing mathematics for teaching within classroom interactions in teacher education. *Educational Studies in Mathematics* 81:1–14. URL http://dx.doi.org/10.1007/s10649-011-9371-9. 10.1007/s10649-011-9371-9.

Shulman, Lee S. 1986. Those who understand: Knowledge growth in teaching. *Educational Researcher* 15(2):4–14.

———. 1987. Knowledge and teaching: Foundations of the new reform. *Harvard Educational Review* 57(1):1–22.

Silverman, Jason, and Patrick W. Thompson. 2008. Toward a framework for the development of mathematical knowledge for teaching. *Journal of Mathematics Teacher Education* 11(6):499–511.

Sleep, Laurie. 2012. The work of steering instruction toward the mathematical point: A decomposition of teaching practice. *American Educational Research Journal* 49(5):935–970. doi:10.3102/0002831212448095. URL http://aer.sagepub.com/content/49/5/935.abstract. http://aer.sagepub.com/content/49/5/935.full.pdf+html.

Sleep, Laurie, and Samuel L. Eskelson. 2012. MKT and curriculum materials are only part of the story: Insights from a lesson on fractions. *Journal of Curriculum Studies* 44(4):536–558.

Speer, Natasha M., and Ole Hald. 2008. *Making the Connection: Research and Teaching in Undergraduate Mathematics Education*, chap. How Do Mathematicians Learn to Teach? Implications from Research on Teachers and Teaching for Graduate Student Professional Development, 303–315. Washington, DC: Mathematical Association of America.

Speer, Natasha M., John P. Smith III, and Aladar Horvath. 2010. Collegiate mathematics teaching: An unexamined practice. *Journal of Mathematical Behavior* 29:99–114.

Speer, Natasha M., and Joseph F. Wagner. 2009. Knowledge needed by a teacher to provide analytic scaffolding during undergraduate mathematics classroom discussions. *Journal for Research in Mathematics Education* 40(5):530–562.

Thompson, Clive. 2011. How Khan Academy is changing the rules of education. online; last viewed Nov. 18, 2012. URL http://www.wired.com/magazine/2011/07/ff_khan/.

Thompson, John R., Warren M. Christensen, and Michael C. Wittmann. 2011. Preparing future teachers to anticipate student difficulties in physics in a graduate-level course in physics, pedagogy, and education research. *Physical Review Special Topics - Physics Education Research* 7:010,108 1–11. doi:10.1103/PhysRevSTPER.7.010108. URL http://link.aps.org/doi/10.1103/PhysRevSTPER.7.010108.

Wagner, Joseph F., Natasha M. Speer, and Bernd Rossa. 2007. Beyond mathematical content knowledge: A mathematician's knowledge needed for teaching an inquiry-oriented differential equations course. *Journal of Mathematical Behavior* 26:247–266.