Spring 5-2020

An Evolutionary Approach to Crowdsourcing Mathematics Education

Spencer Ward

AN EVOLUTIONARY APPROACH TO CROWDSOURCING MATHEMATICS EDUCATION

by

Spencer Ward

A Thesis Submitted in Partial Fulfillment
of the Requirements for a Degree with Honors
(Mathematics)

The Honors College

University of Maine

May 2020

Advisory Committee:
Justin Dimmel, Assistant Professor of Mathematics Education and Instructional Technology, Advisor
François Amar, Dean of the Honors College and Professor of Chemistry
Dylan Dryer, Associate Professor of Composition Studies
Hao Hong, Assistant Professor of Philosophy and CLAS-Honors Preceptor of Philosophy
Michael Wittmann, Professor of Physics

ABSTRACT

By combining ideas from evolutionary biology, epistemology, and philosophy of mind, this thesis attempts to derive a new kind of crowdsourcing that could better leverage people’s collective creativity. Following a theory of knowledge presented by David Deutsch, it is argued that knowledge develops through evolutionary competition that organically emerges from a creative dialogue of trial and error. It is also argued that this model of knowledge satisfies the properties of Douglas Hofstadter’s strange loops, implying that self-reflection is a core feature of knowledge evolution. This mix of theories then is used to analyze several existing strategies of crowdsourcing and knowledge development, allowing the identification of a small number of design mechanisms that combine in different ways to create each strategy’s power. Finally, a website is proposed that combines all of these mechanisms to crowdsource the self-reflective evolutionary development of mathematics education using existing web design techniques.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

INTRODUCTION

Untapped Potential

The Internet is a free-for-all. There is such a high skill curve involved with modern web design that most of the people who end up involved with it are people who have spent most of their lives focused on web design. There is nothing especially wrong with that, but the Internet is one of the most unprecedented tools people have ever discovered. We have been able to make a great deal of progress based mostly on the ideas of those web designers. But that also means that there might be a great deal of potential left entirely untapped. Theories from fields outside the technology world could easily hold promising new ideas, and they may never be discovered.

Spreading Mathematics

Meanwhile, the rest of our knowledge needs the help as much as it ever has. Modern education systems have proven that most people can understand a large variety of important ideas. As much as we should celebrate that fact, the success of modern education highlights a tantalizing new goal still outside our reach. There are plenty of mathematical ideas that could benefit the world in undiscovered ways, but any of those that are not taught in schools are often never learned at all. This is not the fault of the education system, and it is especially not the mathematicians' faults. Mathematical ideas are hard to explain, and we only have so many people trying to explain them, so we are only going to have time to find ways to explain the highest priority ideas.

The Internet is already being used in powerful ways to address that challenge. There are a considerable number of initiatives, which I will go over in Chapter IV, that are using web platforms to crowdsource significant progress in mathematics education.

Evolution and Minds

However, it is possible that evolutionary biology and philosophy have already contributed the ideas necessary to apply the ingredients of those platforms' success in a novel crowdsourcing strategy. Specifically, I will use theoretical physicist David Deutsch's synthesis of Richard Dawkins' theory of meme evolution with Karl Popper's theory of knowledge to understand what allows knowledge to develop. By combining these ideas with Douglas Hofstadter's concept of strange loops, I will attempt to explain how they suggest that a small number of changes to the existing crowdsourcing formulas could yield a new strategy with a new kind of innovative promise.

CHAPTER I:

GENES, MEMES, AND CREATIVITY

Finding a Foundation

Knowledge and crowdsourcing are huge topics. But they do not need to be. Their apparent complexity is largely a result of their conceptual ambiguity. There is no end to the interpretations that could be applied to either, and each of these interpretations has its own set of intuitions and heuristics that could lead to different conclusions. I need to choose my interpretation carefully. They need to be grounded in mechanisms that clearly exist in the real world, but they also need to tell me useful things at both the arbitrarily large scales present in crowdsourcing and the arbitrary levels of depth present in mathematical knowledge.

In this chapter, I will present a modern epistemology proposed by David Deutsch, the inventor of quantum computing (Deutsch, 1985). This will provide a scientific grounding that applies at every scale. In the next chapter, I will use cognitive scientist Douglas Hofstadter's idea of "strange loops" to illustrate how Deutsch's interpretation of knowledge accounts for the level of depth present in mathematics.

Atoms to Genes

Following Deutsch's 2011 book *The Beginning of Infinity*, let us begin with evolutionary theory. Picture a sea of atoms. Since atoms have the potential to bond with each other, some of them might clump into molecules. Molecules can interact with each other, so some configurations of molecules might create chemical chain reactions. It is conceivable, though highly unlikely, that some of these reactions might be able to continue indefinitely. In fact, although even more unlikely than that, it is theoretically possible that some of those indefinite reactions might be able so sophisticated that they make fully functioning copies of themselves. The simplest of these might employ specific mechanisms to replicate specific components of themselves. However, an incredibly unlikely subset of these reactions could, theoretically, unknowingly treat a few of their molecules like a blueprint for reproduction (Deutsch, 2011).

That seems to be what happened with cells. It seems that the sheer size of the universe overcame their improbability. As these cells copied themselves into increasingly large populations, the law of large numbers then guaranteed that at least some of their genes would be miscopied from generation to generation. Those miscopies had a chance to change how the reactions operated, potentially changing the number of times it can copy itself. The populations of cells with traits that led to high levels of copying grew faster than their low-copying cousins, increasing the presence of those high-copying traits in the environment and possibly driving the low-copying cells to extinction. Over time, increasing complexity would have allowed gene miscopies to have increasingly large effects, drastically increasing variation between generations (Deutsch, 2011).

As Charles Darwin predicted in his theory of evolution by natural selection, that variation was enough to transform cells into plants, animals, and people. Like a hacker brute-forcing bank passwords by entering every possible combination, evolution brute-forced survival by randomly branching into huge numbers of mutations. The lines of development that interacted well with events in their environments accumulated far larger populations than their peers. After large amounts of time, the most successful lineages were able to outcompete the others, causing all but a few lines of development to fall away. The only organisms that remained were the ones that fit their environmental constraints (Deutsch, 2011).

Genes to Memes

Fast forward a billion years or so, and a similar effect began to play out between organisms and their behavior. Using newfound abilities to see, hear, and move, they started to interact with each other. Birds had songs, wolves had body language, and apes had gestures. These are, as Richard Dawkins named them in 1976, simple examples of *memes*—behaviors that causally contributed to their spread. Like genes, memes had the potential to change through branching lineages of inheritance with mutation, meaning they developed through an analogous process of evolution by natural selection (Dawkins, 1976).

The earliest memes were likely spread directly by genes. Certain mutations led to social behaviors that increased organisms' chances of reproduction, allowing them and their descendants to spread that behavior (Dawkins, 1976). As time went on, increasingly complex biological mechanisms made those emergent mechanisms more convoluted. It seems that meme evolution outpaced gene evolution, and something else had to take over (Deutsch, 2011).

While most lineages of species were developing ever more sophisticated mechanisms for replicating certain types of memes, one of them apparently stumbled on something far more effective: *creativity*. Deutsch posits that a primitive form of the creative process, composed of (a) making a guess, (b) observing the results, and (c) adjusting the next guess accordingly, would be able to create an effect analogous to gene mutation. Even if an all an organism's guesses began as random neural patterns, any kind of ability to recognize successful guesses would gradually allow it brute-force external causality. An organism with a powerful enough creative organ would be able to transcend biologically encoded memes in exchange for causal models of the world. Deutsch refers to these as *explanations* (Deutsch, 2011).

Memes to Explanations

At some point, the evolutionary value of memes seems to have mixed with the increasing complexity of meme evolution. Genes and memes must have co-evolved, granting apes vocal cords and prefrontal cortices. Those apparently allowed branches of sound-based memes to evolve into spoken language, giving early humans the ability to share rough versions of their explanations between each other. Thus, creativity's approximation of causality began to take place across multiple minds, and the explanations they produced became increasingly accurate at describing abstract reality (Deutsch, 2011).

This, Deutsch argues, is the source of knowledge. As memes (including language) evolve, creativity can operate on increasingly sophisticated guesses and analyses. A growing knowledgebase of language and explanations allows someone to develop their knowledge around ever more abstract sources of causality (Deutsch, 2011).

Explanations to Knowledge

As Deutsch notes, this view of knowledge generalizes the theories of one of the most prominent epistemologists, Karl Popper. As Popper described in his 1963 book *Conjectures and Refutations*, all knowledge must start with a theory. Even if one makes a conclusion from data, Popper argued they must have first produced a theory explaining how that data should be interpreted. Our conclusions can only ever be our best interpretation. However, others can then critique that theory separately from its conclusions, potentially leading to both a better theory and better conclusions. Popper argued that this dialogue of conjecture and refutation produces all of our knowledge (Deutsch, 2011).

Philosopher and mathematician Imre Lakatos extended this theory to mathematical knowledge in his 1976 book *Proofs and Refutations*. He described mathematics as a dialogue of imperfect proofs being improved through refutations. This means that mathematical knowledge evolves through the same creative process as Deutsch's explanations (Lakatos 1976).

Popperian Knowledge Evolution

Rewording both theories in Deutsch's terminology, someone uses language, their existing knowledge, and the first half of the creative process to attempt to present an explanation. Someone (possibly the same person) can then apply the second half of the creative process to compare this attempt's consistency with evidence relative to competing attempts. If they find anything to improve, they can respond to the first person (or themselves) with critique. As this process plays out within individuals and across groups, creativity gradually creates opportunities for explanations to evolve into reliable knowledge (Deutsch, 2011).

This model of knowledge gives us a view into its inner workings. Creativity, explanations, and language haphazardly build on each other to produce knowledge about increasingly sophisticated ideas. The way this theory is framed, this Popperian process is sufficient to explain all forms of knowledge development. That means that any attempt to crowdsource knowledge, including mathematical knowledge, relies on some combination of these same mechanisms.

CHAPTER II:

PERPETUAL ABSTRACTION

Infinite Reach at Finite Speeds

Before I apply Deutsch's theory of knowledge, I would like to discuss its most confusing implication. He named his book *The Beginning of Infinity* because he was arguing that his observations about knowledge imply that it has infinite reach. Given the caveats he makes in the book, that is a fittingly powerful and accurate description of why people seem so special. However, I find that Deutsch's metaphors fail to communicate the weight of his conclusion. They evoke science-fiction style intuitions of grandeur hundreds of years in the future, and Deutsch plays into that by exploring those scenarios. Those intuitions are deeply valuable, but they deserve to be rephrased to portray the full depth of human knowledge.

Because Deutsch's theory of knowledge applies just as well to individuals as it does to groups, it can also operate as a functional theory of mind. That means it can potentially be merged with other theories of mind to produce new implications for epistemology. Douglas Hofstadter's strange loops may be just what I need.

Getting a Grasp on Strange Loops

Strange loops are philosophically slippery. When Hofstadter proposed them in *Gödel, Escher, Bach: An Eternal Golden Braid* (*GEB*) in 1979, he felt that they were such a difficult idea that he spent the majority of the book's 777 pages talking not about minds, but formal mathematics, physics, computers, biology, art, and music. Since then, he has written several other books attempting to get his ideas across in a variety of other ways. I will follow his most recent, *I Am a Strange Loop* released in 2007, but I will focus primarily on the ideas central to *GEB*.

A strange loop is an unbound but time-dependent recursive abstraction system. Strange loops are things that can recurse up and down through arbitrarily many levels of abstraction while never causally depending on anything taking place beyond a base level of abstraction. In *I Am a Strange Loop*, Hofstadter described them as

> an abstract loop in which, in the first series of stages that constitute the cycling-around, there is a shift from one level of abstraction (or structure) to another, which feels like an upwards movement in a hierarchy, and yet somehow the successive "upward" shifts turn out to give rise to a closed cycle. That is, despite one's sense of departing ever further from one's origin, one winds up, to one's shock, exactly where one had started out (Hofstadter 2007, page 102).

Hofstadter introduced his idea using art from M. C. Escher and a pair of mathematical theorems by Kurt Gödel. Escher gives a hint; Gödel gives a proof.

Drawing Hands

In a moment, look over the copy of Escher's *Drawing Hands* on the next page. You may have seen this sketch many times before, but this time, run an experiment. Forget that it is only a drawing and let yourself fall for the illusion for a minute or two. Focus on the specific, subjective experience of moving your gaze from each hand's fingertips to the other hand's wrist. Simply go around the loop until you feel that you understand what you are experiencing, then continue reading.

Figure 1. *Drawing Hands* (scanned from Hofstadter 2007)

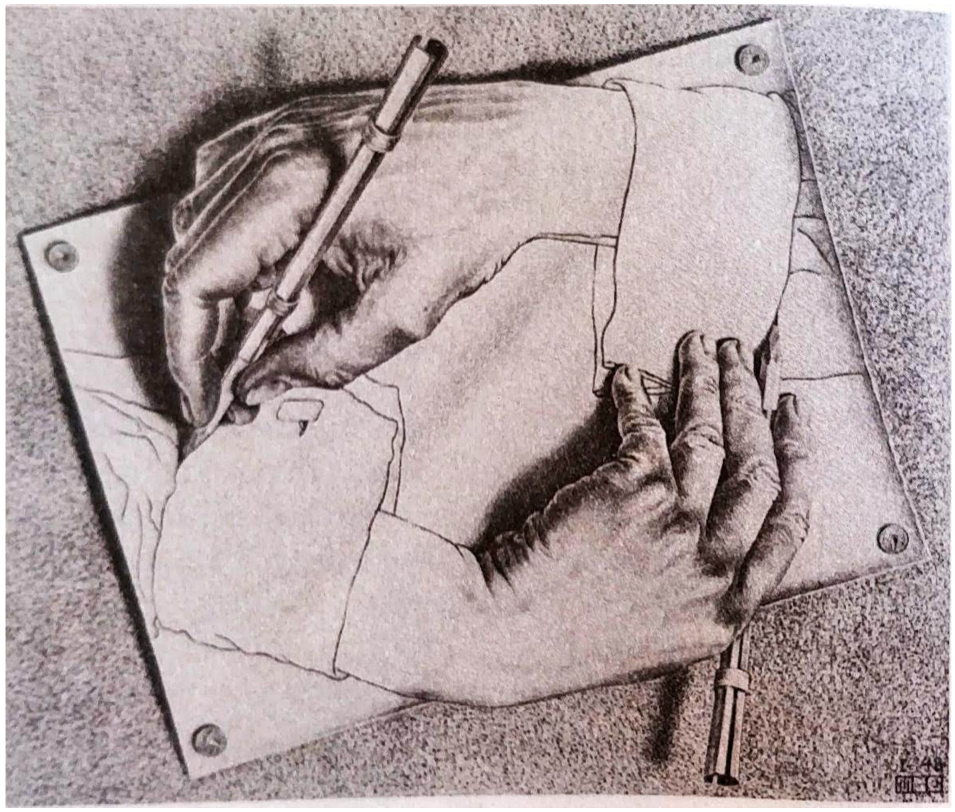

If you were paying close attention, you may have noticed that each hand seems to rise from the basic level of abstraction of a crude sketch to the significantly higher level of abstraction of a lifelike hand floating above the page. When you move your gaze from the realistic hand to the crude sketch it was drawing, the abstraction rises to a new level. Moving your gaze to the new set of fingers lifts abstraction still higher, but it also brings you back to where you started. You can then continue this cycle for as long as you wish, abstracting from hand to paper and back *ad infinitum*.

Much like optical illusions, which try to trick the eyes into seeing something paradoxical, *Drawing Hand* attempts to trick the mind into considering something paradoxical. Optical illusions only work because of specific properties of the subjective experience of vision. Analogously, *Drawing Hands* seems to rely on specific properties of the subjective experience of thought. No matter where you start your gaze, you are always able to abstract through hands and pages to your heart's content. If you wished, you could even abstract from the illusion itself. Every time you look at *Drawing Hands*, you could keep track of the layers of abstraction you recurse through. No matter how big that number becomes, you could always keep abstracting to higher and higher levels of hands. The experience of looking at *Drawing Hands* is evidence that you are a strange loop (Hofstadter, 2007).

Gödel's Comment on Knowledge

Hofstadter realized that this same phenomenon is at work in Kurt Gödel's famous incompleteness theorems (shown below).

Figure 2. Gödel's Incompleteness Theorems. Shown as paraphrased by *The Stanford Encyclopedia of Philosophy*.

First incompleteness theorem
Any consistent formal system F within which a certain amount of elementary arithmetic can be carried out is incomplete; i.e., there are statements of the language of F which can neither be proved nor disproved in F.

Second incompleteness theorem
For any consistent system F within which a certain amount of elementary arithmetic can be carried out, the consistency of F cannot be proved in F itself.

These theorems showed that no formal system can possibly encode the entirety of valid mathematics. Curiously, Gödel found that their incompleteness was not a result of a lack of computational power, but an abundance. Only formal systems that can carry out “a certain amount of elementary arithmetic” are incomplete.

This implies that human-level mathematical knowledge must always be one step ahead of formal systems. In fact, Gödel’s proof is structured around a set of directions. It is true because, in principle, people could always carry out those directions for any new formal system. Thus, it is a statement that the continuing development of human-level mathematical knowledge can always abstract beyond any amount of knowledge encoded in a formal system.

Creative Knowledge Evolution Creates Strange Loops

We can now turn back to David Deutsch. At the beginning of this chapter I mentioned that his theory of knowledge is also into a theory of mind when considered at the level of individuals. It should not be surprising, then, that knowledge evolution is sufficient to create strange loops. As explanations pass between people and are exposed to varying uses of language and creativity, those explanations can evolve to describe situations taking place at increasing levels of abstraction. Since, as Deutsch concluded, this process has no limit, knowledge evolution has no abstraction limit.

This gives an equivalent alternative to Deutsch's conclusion about knowledge's infinite reach. Given any level of knowledge, people are always capable of reflecting on it to learn more abstract things. The knowledge created by a Popperian, creativity-driven evolution of explanations can be arbitrarily self-reflective.

To be clear, this argument does not depend on knowledge teleologically moving towards abstraction. Knowledge produced by creative explanation evolution should be considered as unpredictable and arbitrary as any other products of evolution. However, as long as it is possible for individuals to stumble upon arbitrary ideas that approximately track the world, it should also be possible for others to recognize the validity of those ideas and establish them as part of the collective level of knowledge. Reliable explanations are improbabilities that need to be deliberately identified.

In this light, the way people bring awareness to different parts of their knowledge development process may have a direct effect on their ability to translate improbable individual success to consistent collective success. Explanations evolve regardless of whether people pay attention to them. However, the more attention people do pay, the higher the chance of them finding reliable knowledge. To see how that happens with mathematics education, I will now go over some of the established strategies of developing knowledge.

CHAPTER III:

KNOWLEDGE EVOLUTION IN THE REAL WORLD

Classrooms

Classrooms are a great place to start. Thousands of different educators, each with their own styles built on the styles of people that came before them, are applying their current knowledge to real people and watching exactly how each of those people learn. Educators may need to limit their experimentation somewhat to ensure their students learn everything they need, but their limitation is also balanced by the detailed, immediate feedback that they receive. In the context of the epistemology I have been using, classrooms produce knowledge evolution with small amounts of creative variation between generations. However, the creative reflection and critique that guides evolutionary success is based on direct feedback with real people. That means that the large variety of lines of development all develop in dialogue with real learning.

Academic Research

Academic research helps glue together the knowledge evolution taking place in classrooms while adding a significant amount of self-reflection. The people contributing to it are basing their input on either the valuable classroom experience I mentioned above, or data collected through scientific experiments. This makes the potential creative variation significantly higher than it could ever be in individual classrooms. Academic structure, exemplified in peer review but present throughout the academic world, counters the stream of creative input with highly abstract standards based on many years of reflection on the way knowledge has developed in the past.

The Power of the Internet

Recently, the Internet has become another major source of knowledge development. Content hosted almost anywhere online can reach just about anyone on Earth within a second or two. Better yet, everyone sees the same content at the same time; it is globally synchronized. And unlike all the other telecommunications technologies that led up to it, the Internet can be used for just about anything. Modern computers, even if we only consider their web browsers, support immense amounts of design flexibility. User interfaces can be used to create effortless mechanisms that would have been unusable on pen and paper—imagine Facebook if people had to submit their posts through mail. These features make the Internet perfect for education and crowdsourcing. And they have been used to great effect.

Private Companies

Consider sites like *Khan Academy*, *Brilliant*, and Wolfram's *MathWorld*. Each takes a completely different approach to presenting educational content. They build on different pools of pedagogical knowledge, different groups of educators, and different audiences of learners. However, each of these sites allows the knowledge of those individuals to accumulate over time. Sometimes, parts of that knowledge can spread to the competitors. Other times, they each maintain their own line of development. Occasionally, new sites arise to fill in gaps in the competition. In effect, then, they collectively feed knowledge evolution. Their contributions are haphazard, but their content is shared on such a large scale that its messiness is a non-issue.

Wikis

A cleaner route to a similar outcome could be wiki-style crowdsourcing. As popularized by Wikipedia, the wiki model is based around giving every idea its own page. Users are then allowed to edit the page's content to produce new versions. Others can validate those changes and improve them as they like, providing the effects of an attempt/critique dialogue with the critique left implicit. Versions develop in a linear history, but they still do evolve. The sheer scale of people taking part ensures that a page's linear development accumulates many potential insights produced by other explanations evolving outside of the page, creating a reliable stream of knowledge development.

Questions and Answers

Question and answer (Q & A) websites like *Stack Exchange* and *Quora* harness a similar but more responsive form of knowledge evolution. Content develops through a direct dialogue between educators and learners. Learner questions offer something akin to critique of widely available mathematics education. Any educator who sees the question can respond with creative attempts to address those critiques with answers, producing a competing pool of explanations. Users can vote on question quality, attracting attempts to the pressing critiques. Since people can also vote on answer quality, these high-quality questions are then paired with their highest quality answer. This kind of crowdsourcing allows knowledge to evolve based on need, directly decided by the people watching it evolve.

Each of these strategies offers unique and powerful benefits for knowledge development. As I have tried to illustrate, each one creates its power by supporting different mechanisms of knowledge evolution. However, I see no reason why any of them would be mutually exclusive. If knowledge really does develop through the combination of all these mechanisms interacting as a whole, then it might be possible to design a crowdsourcing site that supports them all.

CHAPTER IV:

DESIGNING AROUND KNOWLEDGE EVOLUTION

Partial Coverage

This is where theory starts to pay off. If it is indeed true that all knowledge develops through creatively derived explanations evolving as memes, then all the mechanisms I described in chapters two and three are at least implicitly at work in the strategies I discussed in the last chapter. However, each strategy seems to be based on a different understanding of the way knowledge develops. That means that each one explicitly emphasizes different parts of the evolutionary process. See Table 1 for a rough idea which mechanisms each prioritizes (note that the development process is a wildcard for private companies).

Table 1. Combinations of Evolutionary Mechanisms

	Creative Dialogue: Agency	**Creative Dialogue: Awareness**	**Inheritance with Variance**	**Competitive Selection**	**Reflection on Development**
Classrooms	Explicit	Explicit	Implicit	Implicit	Implicit
Academic Research	Explicit	Explicit	Implicit	Explicit	Explicit
Private Companies	?	?	Implicit	Implicit	?
Wikipedia	Explicit	Implicit	Explicit	Implicit	Explicit
Q & A	Explicit	Explicit	Implicit	Explicit	Explicit

Every strategy produces knowledge no matter which mechanisms it leaves implicit. Knowledge can evolve fine without any planning. However, when people pay attention to specific evolutionary mechanisms, they can apply creativity and produce explanations to make those mechanisms an object of their knowledge. That opens the potential for people to improve the way those mechanisms operate, leading to more effective knowledge development.

Full Coverage

Applying this lesson to crowdsourcing means explicitly designing around those mechanisms. Specifically, a hypothetical full-coverage website would need to accomplish the following:

1. Separate creativity into two separate steps
2. Allow content to develop through a system of inheritance with variation
3. Track competition between lineages of development
4. Allow users to understand and moderate the evolutionary process

Paradoxically, this may be far easier than it sounds. The first three goals are straightforward—tricky, but nothing outside of normal website design. And once a website fulfills those, it will contain the full detail of the Popperian evolution that Deutsch argues makes up the entirety of knowledge development. That means that the fourth goal becomes a much more palatable matter of applying standard design tools to crowdsource feedback on the mechanisms behind the first three goals.

I will return to these design goals in Chapter VI. First, however, I need to define what kind of content this kind of website should have.

CHAPTER V:

MATHEMATICS FOR WIDE AUDIENCES

Simple by Design

So far as I can tell, simply adding the features I just described to a traditional wiki would already be a valuable experiment. That seems like a bit of a waste, though. Wiki-style content has certainly proven its worth. But a lot has changed since the wikis were invented.

When Ward Cunningham created the wiki format with *WikiWikiWeb* in 1995, he did not want to use any special design techniques or content development processes (see http://wiki.c2.com/?WhyWikiWorks for *WikiWikiWeb's* view of itself). Computers and Internet connections were several orders of magnitude slower than they are today, severely limiting the design potential of web content. Simplicity was a huge design priority for everyone. That was one of the reasons why wikis were so important. They were the simplest way for people to share knowledge on a large scale, and their communities celebrated that fact (see http://wiki.c2.com/?DoTheSimplestThingThatCouldPossiblyWork). However, twenty-five years later, even the slowest computers and Internet connections support far more design complexity than is present in wikis.

This is useful to note, because I am not sure the wiki content format would be flexible enough to support my overarching goal of spreading mathematical knowledge to as many people as possible. Content needs to support audiences with wildly different levels of background knowledge. Thanks to the design flexibility provided by modern technology, I might be able to find new techniques that better support those audiences while preserving a simple user experience.

Connections and Metaphors

I can start by determining what it means, within the theory of knowledge I have been using, for people to have different levels of background knowledge. If everyone's knowledge is built up through creative explanation evolution, then each individual’s knowledge could theoretically be represented with a unique tree of explanation development. The people with more knowledge on a given topic have somehow used their creativity to stumble upon lineages of explanations that better match the real world. In a very abstract sense, then, adjusting for varying levels of background knowledge means finding a way to help people apply their creativity to develop more accurate explanations. To make that more concrete, I will need a dose of cognitive science.

Imagine two early animals beginning to develop creativity and explanations. The first animal considers each new phenomenon as a brand-new source of causality, and thus it develops every new explanation as a unique cognitive mechanism. The second animal has a mutation that allows it to compare new phenomena to sources of causality it has already explained, allowing it to recycle cognitive mechanisms for new concepts. Both creatures are going to need to approach new ideas through the creative process. They will both come up with guesses, try them out, and observe the results. However, where the first animal will need to start with primitive guesses and abstract up to each new concept, the second would be able to start with guesses as abstract as any of their existing concepts. The first would need to apply its creativity to develop increasingly abstract explanations until it reached the general idea of a concept, then it would need to fine-tune its cognitive mechanisms with observations. The second would jump straight to fine-tuning existing cognitive mechanisms. If the second creature is evolutionarily possible, it would have a huge cognitive advantage over the first. Its descendants would win any evolutionary competition that even remotely involves intelligence (Lakoff and Johnson, 1999).

Not surprisingly, people seem to be descendants of that second creature. As linguist George Lakoff and philosopher Mark Johnson argued in their book *Philosophy in the Flesh* in 1999, brains appear to link neural structures delegated to some concepts with neural structures designed for others. Through mechanisms Lakoff and Johnson call conceptual metaphors and conceptual blends, simple cognitive structures are linked together to perform the work of more complex ones (Lakoff and Johnson, 1999).

In their 2000 book *Where Mathematics Comes From*, Lakoff and psychologist Rafael E. Núñez illustrated how such a system would be able to create mathematical knowledge. As they discuss, only the numbers one, two, three, and possibly four seem to be built into human biology. Newborn babies can distinguish between and accurately perform addition and subtraction within those three or four numbers, but that is all. Lakoff and Núñez propose that all other mathematical knowledge is formed through embodied experience. The causal inputs from simple experiences such as movement, object collection, and interactions between objects always work the same in every situation, so they are a perfect target for creative explanation. These foundational tools, which Lakoff and Núñez label image schemas, are then available as intuitions people can use to accurately make predictions about the world. Mathematics is constructed by structuring those intuitions into more abstract concepts, allowing people to reason about numbers through metaphors to experience. The underlying intuitions produce consistent and precise results because they have been trained with a consistent and precise world. Lakoff and Núñez argue that this is why mathematics feels so real. The symbols and abstractions behind it are purely human creations, but they are making statements about real phenomena (Lakoff and Núñez, 2000).

These insights strongly suggest that conceptual connections and real-world metaphors are foundational to personal knowledge evolution. If I want my design to support wide audiences, those two tools should be central features of any piece of content.

Making Metaphors Concrete

This is important because mathematical ideas are typically not clear from their official statements. On one hand, professional mathematicians mostly deal with math composed of formal definitions, theorems, and proofs that mean specific things. Ideally, people learning about math should be exposed to those formalisms as directly as possible so they can get a concrete reference point for how professionals see their discipline. On the other hand, formalism can often be cryptic to the point of being meaningless. Metaphorical language helps, but maybe not as much as the evidence above would suggest. Metaphorical language can easily drift away from the structure of formalisms. Think of calculus students who understand that a derivative is a curve's rate of change yet have no clue why its definition talks about continuity and limits. Therefore, I would like my design to explicitly connect metaphorical language to the structure of mathematical formalisms. That is to say, I would like to map specific segments of each concept's metaphorical language to specific segments of its formalism. I have sketched one strategy for this below. When the user's cursor hovers over a specific part of a metaphorical explanation, the site would highlight both that part of the metaphor and the corresponding part of the formalism.

Figure 3. Metaphor-formalism Mapping. Different colors signify corresponding sections. On the real site, these would be highlighted individually as the user hovers over them.

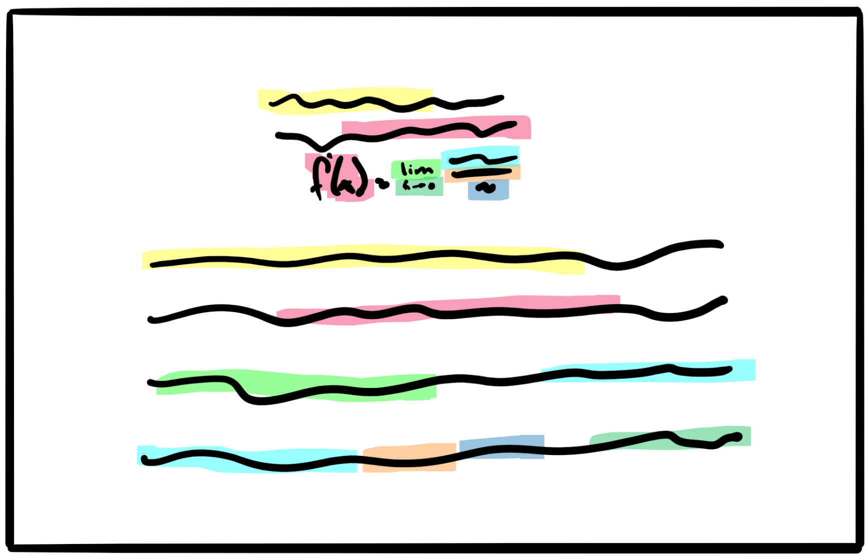

<u>Definitions as Connections</u>

To address the matter of connections, content could be automatically linked together based on its structure. This could be accomplished by choosing a category of mathematical ideas and structuring all the other categories' ideas according to their connections to the central ones. A clear candidate for a central category would be definitions. After all, nothing in mathematics can exist without them. They are what allow us to talk about anything and everything. Every definition, then, would be guaranteed its own page. Anyone who adds a concept that uses that definition would have the choice to give non-definition pages their own pages, but all concepts would automatically populate the pages of any definitions they use. If a concept refers to a definition that does not exist yet, its page would be automatically generated and populated with that concept. That process will become more concrete as I explain more about these pages.

Since mathematical concepts are about more than their formal statement, pages should also be able to include context and discussion about the concept. This could be organized similarly to the Wiki layout, with user-defined sections and an automatically populated table of contents.

Inline Expansion

Ideally, conceptual connections should be a big part of those sections. However, all that extra detail could easily bog down the readers with more background knowledge, making it harder for them to find the information they do not know. Wikis try to solve this to some extent by embedding article links at the first mention of every concept. Wikipedia's latest implementation of this is especially effective, giving web users the ability to hover over every link to read a floating preview as seen in Figure 4.

Figure 4. Wikipedia's Floating Previews

However, this kind of solution struggles beyond basic references. For instance, there is no easy way to get more details on implicit appeal to properties of the reals. Even an explicit reference assumes a certain level of background knowledge. Those details sometimes have their own page, but the floating preview usually cannot provide nearly enough detail to explain something properly. That means the user is forced to set aside the topic they are learning and move to a whole new page. If people were abstract computers sucking up knowledge, maybe that would be fine. However, at least in my experience, that process is both distracting and disruptive.

One way to deal with this might be something I will call "inline expansion." When an educator mentions a definition or a concept that already has a page, the editing page could automatically recognize that as a reference. If a learner wants to learn more about that concept, they could click on the link in the same way they would on a wiki. Instead of opening a floating preview or a new tab, the referenced page would expand between the concept's line and the next. The learner could then read about this concept alongside its usage, making it far easier to apply in context.

Similarly, when an educator references a line of reasoning or an implicit detail that does not need its own page, they would have the option of adding that as an inline expansion. The learner would then be able to click on a corresponding word or phrase to get more detail as they read. Conceivably, both kinds of inline expansions could include nested expansions, which would provide additional granularity for especially complex concepts. Figure 5 gives an idea how this would look. This same system could then be ported to formalisms, allowing users to read about unfamiliar concepts or tricky steps as they make sense of the math. Figure 6 gives a sense how that could work.

Figure 5. Expanded Inline Expansion. Clicking the blue areas would collapse the expansion. Clicking the question mark would signal that something about this page is confusing. Blue arrows signify that the page is scrollable.

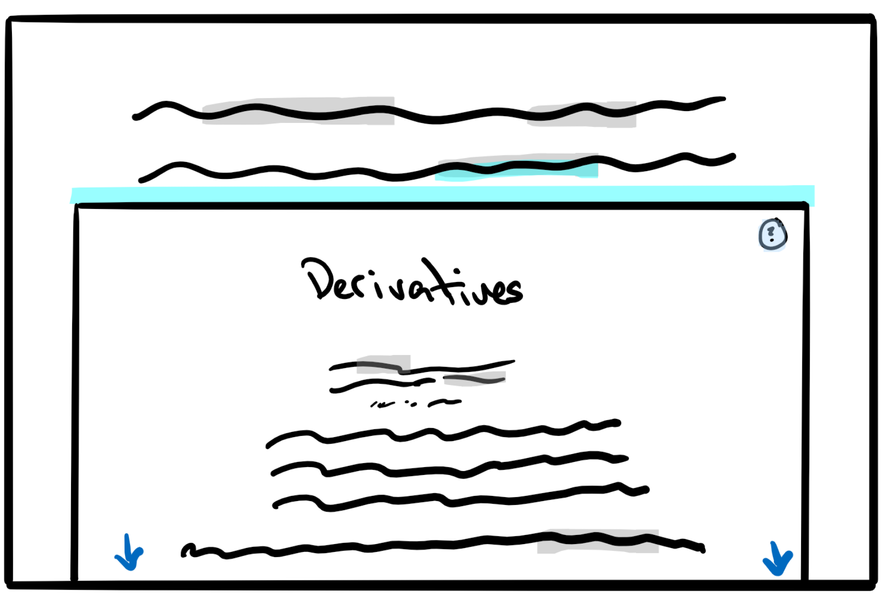

Figure 6. Inline Expansion in a Definition

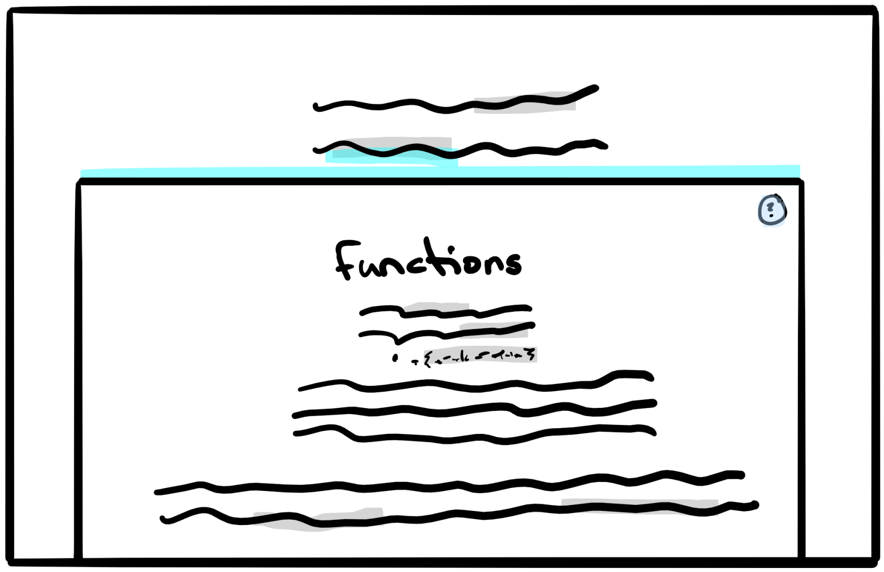

<u>Connected Concepts</u>

Inline expansions will be especially helpful for the last section of each page, the collection of concepts connected to the topic. This is what I was referencing when I mentioned that pages would automatically populate with their related concepts. Each page would include a list of formal usages, such as theorems or proofs that the page's concept is used in. This would be complemented by a second list of the more abstract usages, such as applications, relevant situations, and more abstract ideas (like "Set Theory" or "Linear Algebra"). For the sake of usability, each of these lists will be collapsed by default, as in Figure 7. They can be clicked to be expanded as seen in Figure 8.

Figure 7. Collapsed Concept List. Clicking inside either list would expand it.

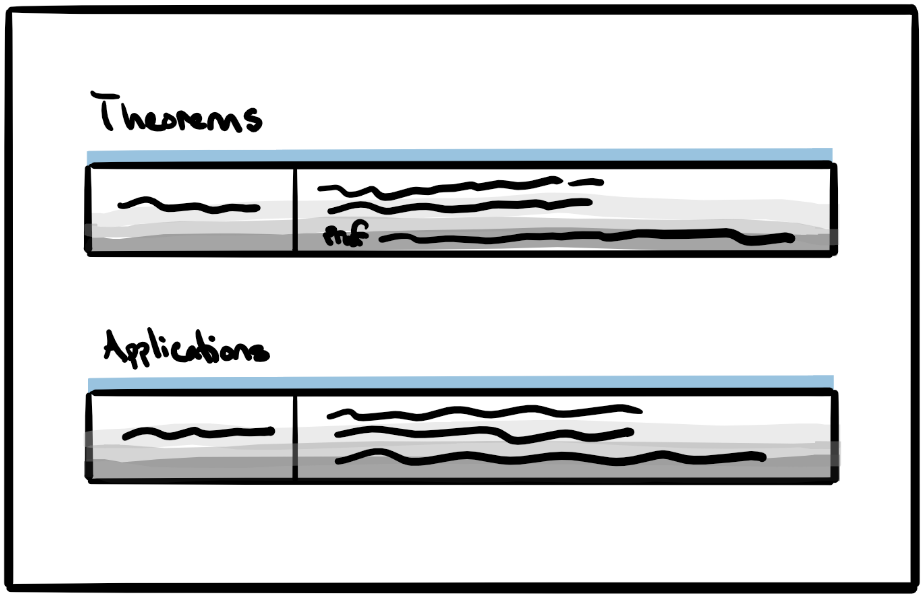

Figure 8. Expanded Concept List. Clicking the blue bar at the top would collapse the list. Users can open theorems on the left and browse their pages on the right. Clicking on the name of a theorem that is already open would open its full page in a new tab.

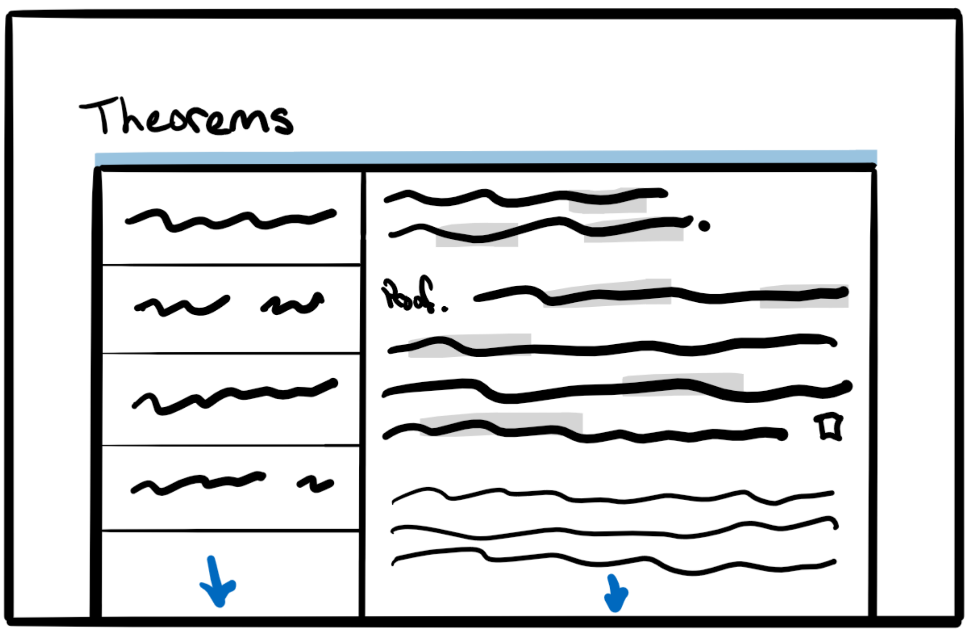

Users could then browse these subtopics as they read about the overarching concept. If the subtopics have their own pages, the user could open them in a new tab by clicking on the topic's title (since that kind of disruption is much more acceptable when someone deliberately wants to learn about a new topic). As with inline expansions, the user could collapse each list by clicking on its top bar.

Since I do not want to limit layouts to the arrangement I would imagine, I intend these three tools of metaphor-formalism mapping, inline expansion, and usage lists to be used as flexibly as possible. For that reason, I will leave their exact ordering up to the editors. For what it is worth, however, I would imagine most pages would look something like Figure 9.

Figure 9. Sample Page Layout. In this case, the definition comes first, then the metaphorical explanation, then context, then related concepts.

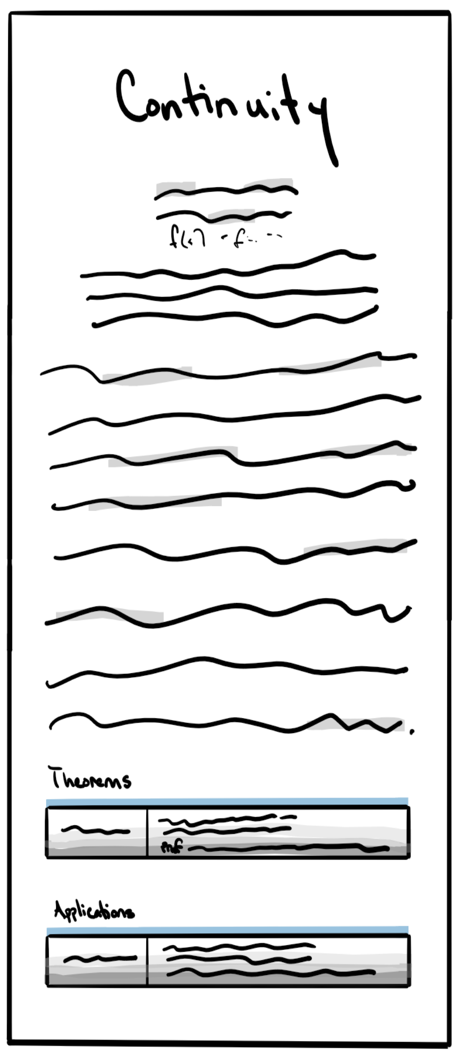

Gathering Feedback

Taken together, these new additions will make content interactive and flexible. Interestingly, this new interaction introduces a possibility that wikis do not have. Say for instance that, for each learner, the site kept track of (a) how much time they spent on each metaphor-formalism mapping, (b) which inline expansions they used, and (c) which expansions they reported as confusing. This data could then be aggregated and shown to educators, allowing them to see how their submissions perform with real world learners. I will discuss this more in the next chapter, but these sources of feedback would allow the awareness and reflection taking place among educators to include data from real learners. This data will certainly be far less detailed than anything produced by academic research or classrooms, but it will still provide more nuance than is available with *Wikipedia*, *Stack Exchange*, or *Quora*.

To help identify where confusion is the most intense, these structured data streams could also be supplemented by instructing learners to click on difficult passages. This may sound like a crude form of feedback. The data will be filled with noise from random clicks. However, confusion should differ far less from person to person than accidental clicks, so the signal may drown out the noise with a large enough data set.

Of course, any kind of data collection raises ethical concerns. Once these design decisions are more concrete, those concerns should be considered as explicitly as possible. In the meantime, I can suggest a few guidelines. Users should be given a clear introduction to how this data will be used, how it will be anonymized to protect their privacy, and an obvious way to opt-out. There is no reason why this site would need to track anything other than the four data sources I have mentioned. These are simple enough that they can be explained to users quite easily, meaning that people can be fully informed of what kind of information they are providing. In the off chance that there is something about this kind of data collection that is impossible to do while fully respecting users' rights and privacy, then these features can be abandoned entirely. Educators should still be able to judge content even without real data.

With that in mind, we can now turn to the half of the design that will produce that content.

CHAPTER VI:

IMPLEMENTING KNOWLEDGE EVOLUTION

Design Goals

To review my goals from Chapter IV, I want the content development side of my website to do the following:

1. Split creativity into two separate steps
2. Allow content to develop through a system of inheritance with variance
3. Track competition between lineages of development
4. Allow users to understand and moderate the evolutionary process

Splitting Creativity

To humor Popper, I will start with design goal (1). Every page will need to start with user submissions. To keep competition high, pages could be open to as many new submissions as people want to write. Users could then rate each submission, and the one at the top of the ranking could become the established public page. This will correspond to Popper's 'conjecture' step.

Each submission could be left open for other users to reply with critiques. These could serve as the primary platform for users to point out problems that need to be addressed in any submission. Hopefully, these will give contributors a chance to improve their work by hearing others' perspectives. These will correspond to Popper's 'refutation' step.

To complete the Popperian dialogue, users could respond to one or more critiques with improvements on the submission being critiqued. To make sure critiques are still addressed, these improvements would inherit all its parent's critiques that it did not address. In fact, these improvements could, themselves, be fully fledged submissions, meaning they could be rated and ranked just like their standalone cousins. Thus, content can develop as Popper observed knowledge naturally develops, through an ongoing dialogue between critiques and improvements. Since the two steps are separated, the website can gather input from people who only want to submit original content, people who can give feedback on that content, people who can follow other people's suggestions to improve existing content, and people who want to both critique and improve existing work. This could be organized as shown in Figure 10 with a submission above and critiques, or with the two swapped.

Figure 10. Sample Submission Page. The links on the left could be used to go to the parent submissions or critiques. The "Critique" button allows the user to submit a new critique. Clicking any of the critiques below would open that critique above, with its corresponding improvements below

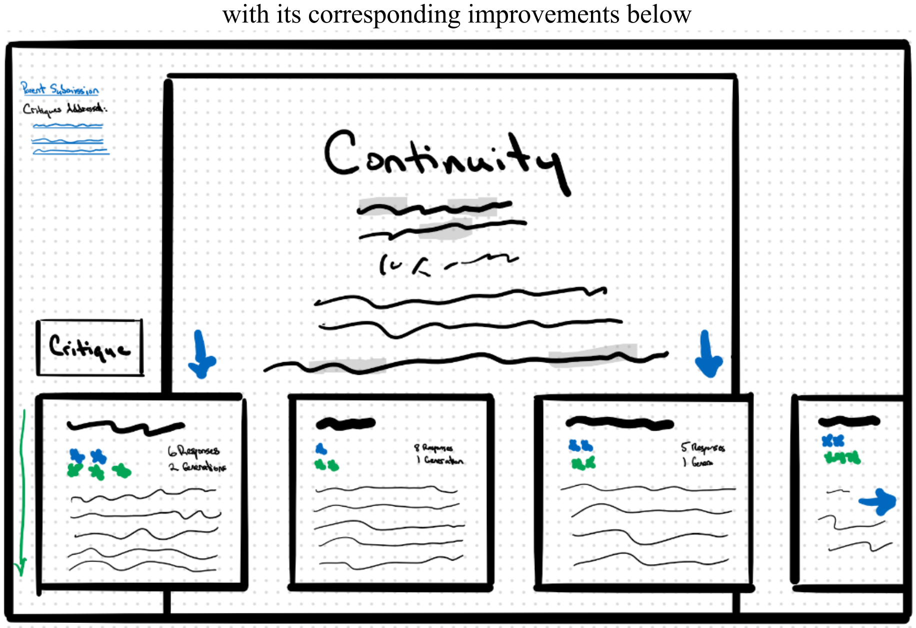

Evolution for Free

Because of the connection Deutsch painted between Popper and Dawkins, the simple implementation I described above gives me design goals (2), (3), and part of (4) for free. Improvements responding to critiques inherit most of the parent submission's content with some variance, accomplishing (2). Since each improvement from a given submission has the potential to branch off into its own line of developing submissions, and each one of those will be competing together using the rating system I mentioned, my design also accomplishes (3). Each standalone submission would therefore have the potential to grow into its own lineage of developments, allowing the website to take advantage of evolutionary principles while scaling to crowdsourcing.

And, as a consequence of the abstraction I described with Hofstadter's theory, the fact that this evolutionary development is being judged by users with the ability to explore all of the different lines of development means that evolutionary success will also be moderated by people's understanding of evolutionary development, accomplishing goal (4).

It is worth noting that, although this style of browsing would take time to use, that may be a good thing. Contributors would gain a significant amount of context by following content's evolutionary development before they rate, critique, or improve it. If users could bookmark certain submissions or critiques, they could find their way back to important places without needing to climb through the whole tree every time. It may be possible to speed that up with a visual navigation system. However, I am not skeptical about how useful such a system would be with large trees. The unpredictable balance between horizontal development (many attempts in one generation) and vertical development (many generations) that needs to be managed in every branch of every lineage could very well make it impossible or impractical to visualize them in a web format. I have not included visual navigation in my current design for these reasons. Instead, I am hoping that the inconvenience will become a natural motivation for ambitious contributors to import the most valuable parts of their favorite lineages into brand new lineages of their own. If they do so successfully, they have a chance of attracting anyone else who wants to move on from the older, more convoluted lineages. That gives each page a natural defense mechanism against excess complexity.

Quality and Potential

The feedback system I mentioned in the last chapter lets me play a trick with the rating system. If the version shown to the public is collecting feedback, that could show up in a usage heatmap as shown in Figure 10. As long as users are asked to rate for quality, other top-rated submissions could also be shown to the public. Using a technique known as A/B testing, different versions could be shown to different users, allowing multiple submissions to receive feedback at once. I will discuss a rating system in a moment, but assume we have one. The highest ranked submission could be shown to the public all the time, but the next best could also be A/B tested alongside it. Once the next best has accumulated enough data, it could be removed from the main ranking and re-rated by the community. If it rates higher than the primary version, it would replace it; otherwise, it could be listed in a secondary ranking exclusive to publicly tested submissions.

Since the ranking of high-quality submissions is gradually being trimmed from the top, it is actually somewhat desirable to use a rating system that skews towards higher rankings. One such system, which would preserve both granularity and objectivity would be a three-star rating system. Users could be instructed to use one star to designate anything from poor to mediocre; two stars could represent anything acceptable but needing improvement, and three stars could represent anything worth showing to the public. Much of the granularity separating, say, a ‘7/10’ from a ‘10/10’ would be lost. However, since anything rated high enough (the exact cut-off point would need to be determined through user testing) would eventually be A/B tested for comprehension, educators will receive far more nuance in return.

To encourage users to reflect on not just quality, but evolutionary development, they could ask to rate each submission by potential. One star could be designated for no evolutionary potential, two could be designated for an average amount, and three could be designated for anything that deserves to be improved. Submissions with the highest potential could then be ranked alongside the quality ranking, as I will discuss shortly. These would allow educators to identify promising ideas regardless of the quality of their current implementation. If one of those ideas' descendants either achieves a certain quality rating (which would need to be determined through user testing) or gets to the point of being A/B tested, the ancestor could be trimmed from the list. As with the quality rating system, this would trade a small bit of initial granularity for a large amount of long-term nuance. The resulting submission page is sketched in Figure 11.

Figure 11. Sample Submission Page with Feedback and Rating. Areas of the page that people click the most are highlighted in darker red. The stars on the right would be used to select a rating. This user has chosen to rate this page 'one' for quality and 'two' for potential.

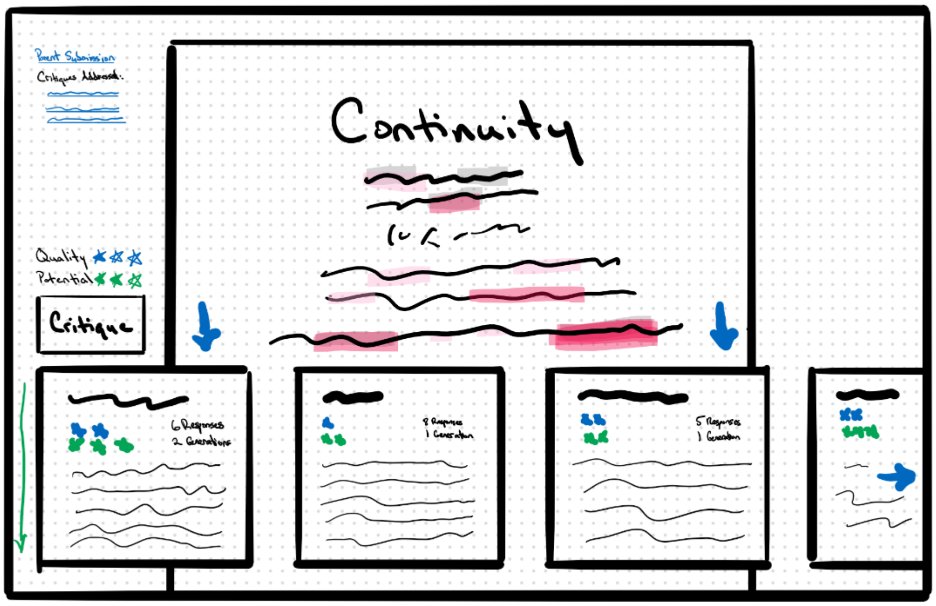

Editing Hubs

To bring this all together into something concrete, the main editing hub of every page could look something like Figure 12. At the top of the page, educators would see a list of the top-performing submissions that have been A/B tested (including the version the public sees). Below that, they would see a list of the top-performing submissions that have not yet been tested. Since performance is measured by quality and potential separately, these lists could be sorted by either. I have chosen to let users toggle between the two, with potential set as the default. However, it may turn out that separating the quality and potential sortings of tested and untested submissions into four separate lists leads to better development. It may also be possible that quality and potential should be averaged or multiplied together to make a single ranking. These options should be explored during user testing.

Figure 12. Sample Editing Hub. Currently, the two rankings are sorted by quality and potential, respectively; each of these could be toggled by clicking on the highlighted words 'Quality and Potential.' Clicking any of the headers in the 'Original Submissions' list would sort the list by that measure.

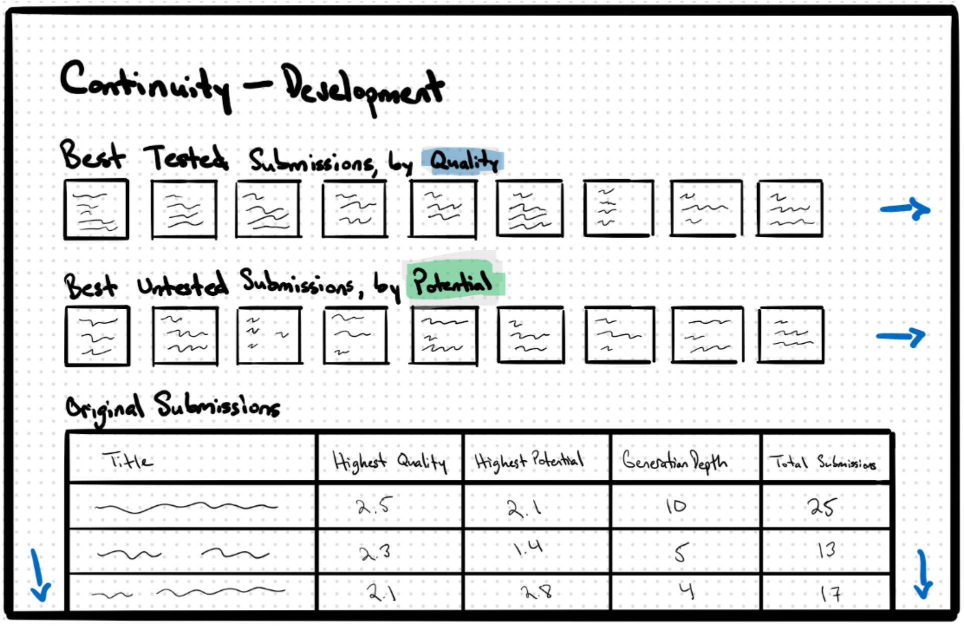

Below those two (or four) lists, educators could be shown a table of all the topics' base submissions—standalone submissions forming their own lineages. As seen in Figure 12, each lineage could be listed with (a) the highest quality rating of any submission it contains, (b) the highest potential rating potential rating, (c) the largest number of generations of critique and submission in any of its branches, and (d) its raw number of submissions. Since this is a table, users could choose to sort by any of these measures. As with the other two lists, though, I would make potential the default.

As I have described it, this edit screen provides a route into each lineage along with shortcuts into the most important submissions in any lineage. This gives users a way to jump straight into important lines of development, pointing crowdsourced effort wherever it is needed most.

Taken as a whole, this design makes all the mechanisms of knowledge evolution into explicit design elements. It separates the creative process into a traceable dialogue. It allows different ideas to branch off while keeping the important parts of existing work. And it lets people moderate content evolution based on their abstract understanding of it. Combined with the user feedback and content flexibility introduced in the last chapter, this design should produce a unique kind of innovation in math education.

CONCLUSION

It is tough to say whether it is harder to wrap one's head around the simplicity of evolutionary principles or the complexity of their implications. Luckily, that oddity means that I only need a modicum of understanding of either in order to exploit the simplicity to address the complexity. As I tried to show in Chapters V and VI, modern web design tools seem perfectly capable of implementing the kind of features it would take to crowdsource a self-reflective evolutionary dialogue of knowledge. This could be used to address the enormous challenge of spreading complex mathematics, but in principle it could also be applied to just about anything.

Of course, the design presented in this thesis is only an early attempt at applying these ideas. There is a considerable amount of work left to be done in refining my use of philosophy to the level of a trustworthy foundation. The features and designs I have presented are early applications of that early theory, so they should be judged with even more skepticism.

The full platform will need to be refined, prototyped, and tested. The philosophy I have applied so far will need to be supplemented by evidence from psychological and pedagogical research. Branching decision trees will need to be sought out and explored. It will be complex, for sure. But it might just work.

BIBLIOGRAPHY

Dawkins, Richard. *The Selfish Gene*. Oxford University Press, 1976.

Deutsch, David. *The Beginning of Infinity*. Penguin Group, 2011.

---. "Quantum theory, the Church–Turing principle and the universal quantum computer." Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences 400 (1985): 117 - 97.

Hofstadter, Douglas. *Gödel, Escher, Bach: An Eternal Golden Braid*. Basic Books, 1979.

---. *I Am a Strange Loop*. Basic Books, 2007.

Kant, Immanuel. "What Is Enlightenment?" Columbia University, www.columbia.edu/acis/ets/CCREAD/etscc/kant.html, 1784

Lakatos, Imre. *Proofs and Refutations*. Cambridge University Press, 1976.

Lakoff, George and Mark Johnson. *Philosophy in the Flesh*. Basic Books, 1999.

Lakoff, George and Rafael. E. Núñez. *Where Mathematics Comes From.* Basic Books, 2000.

Popper, Karl. *Conjectures and Refutations*. Routledge, 1963.

Raatikainen, Panu. "Gödel's Incompleteness Theorems." *The Stanford Encyclopedia of Philosophy*, Stanford University, 2 Apr. 2020, plato.stanford.edu/entries/goedelincompleteness/.

AUTHOR'S BIOGRAPHY

Spencer Ward was born in Lewiston, Maine in 1998. He graduated from the Maine School of Science and Mathematics in 2016, leading him to attend the University of Maine with a major in mathematics and a minor in computer science. He is a member of the Alpha Tau Omega fraternity, a former developer, designer, and project manager at ASAP Media Services, and a recipient of an Undergraduate Research and Creative Activity Fellowship from the McGillicuddy Humanities Center.

Upon graduation, Spencer plans to pursue a graduate degree in education research, hopefully leading to a career studying philosophically and scientifically grounded applications of technology.