A summary of my work on concepts

This is a bit long, but broken up into nice short sections. My typographical conventions are to use **boldface** to introduce new terms, and to use *italics* when discussing concepts and properties. Often I will enclose a term in single quotes when I mean 'what is called' .

Preface (for the mathematical audience)

These thoughts arose as I was reading Torkel Franzén's book *Gödel's Theorem: An Incomplete Guide to Its Use and Abuse* . I still haven't read it fully, but I like the idea of it very much, and I'd even recommend it. (Postscript, I realized that another contributing factor was a series of emails written in May 2006 about how to understand context and type information in the calculational style. There I urged that there were only symbols and their contextual properties, and indeed this is all that is left in my conceptual interface for understanding concepts.)

Shortly after becoming involved with calculational mathematics, I began to identify mathematics with various axiomatic systems, and took to calling mathematics "completely precise" . After reading just a few pages of Franzén's book, I was (thankfully!) snapped out of this dream. At one point Franzén is talking about axiomatic systems, and says something like: "If I say that Goldbach's conjecture is true, I don't mean that it is provable in some formal system, I mean that every even number greater than 2 is the sum of two primes." . I realized that he was right on the mark: concepts like *prime* , *even* , *number* , and *2* are not completely precise, even though we sometimes use their completely precise correlates to study them. They are human concepts, just like *dog* , *love* , and *taco* . The former concepts are called 'mathematical' , while the latter are not, but they are all human, and all hazy, to some degree.

These observations led Apurva Mehta and I to discuss concepts and mathematics, and the end result was the following simple interface.

Concepts and Properties

In our minds we have certain psychological objects that we call **concepts** . It is fruitful to think of concepts as clusters of **properties** or attributes. For instance, the concept *apple* has associated with it the properties *crunchy and sweet* , *grows on trees* , etc. These clusters are not necessarily psychologically precise, in any sense of the word. We may have absolutely no conception of all the properties associated with a concept, and the properties we do know, we may not be able to articulate clearly.

Another fruitful way to think about properties is as links between concepts. Under this view, the concept of *apple* would be linked to the concepts of *crunchy* , *sweet* , *growth* , *tree* , and so forth. The properties of *apple* are then just the concepts it is linked to. I like to take a blend of both approaches: I believe that 'clusters' and 'connections' are both psychologically very real.

Please note too that the mind forms concepts, subconsciously or consciously: Our minds have the ability to organize perceptions into new concepts. Unsurprisingly, subconsciously-formed concepts are usually hazier than consciously-formed ones.

Concepts and Names

Before we go on, I want to clarify a point. Earlier I mentioned 'the concept *apple* '. Some people might wonder whether I'm talking about the word 'apple', or the psychological notion of 'apple'. I definitely mean the latter, though since I'm carrying out this discussion through language, I have to use the word 'apple' to invoke that concept. My apologies for any confusion.

Since humans communicate with each other (and themselves), and since psychological reality overlaps among humans to a great degree, we give names to concepts. Names are crucial tools in the reasoning process, but they are only aids to discussion and contemplation, not essential parts of concepts. Different cultures may give different names to the same concept. A baby understands many concepts before it learns its culture's system of names. And so on.

Mathematical concepts

Now, there is a class of concepts that we call **mathematical**. This class is imprecise, but psychologically real. (Just like the class of things we call 'tall', for instance.) Some things are definitely mathematical —like algebras—, other things are definitely not —like love—, and other things straddle the border —like number—. Because the class is imprecise, any description of it can at best be suggestive. But I like to describe it by saying that mathematical concepts have, roughly speaking, fewer properties, and properties that are more precise than those of non-mathematical concepts.

The fewer properties a concept has, and the more easily we can articulate its properties, the more mathematical that concept may be said to be.

Interfaces and Abstraction

Above I mentioned that we use names to help us reason about concepts, because we cannot discuss the concepts directly. This process of giving names to concepts is one example of forming an **interface**. Recall that an interface, in its common usage, is just the medium through which two things interact. For example, a debit card is an interface between us and our money. Apparently we use such interfaces because it is easier to deal with the interface than it is to deal with the raw materials on the other side. (In the above example, money.) There are many examples of interfaces in the world, and especially in the domain of reasoning.

Some very helpful interfaces are formed by a process called **abstraction**. Most concepts have masses and masses of properties, and most of these properties take the form of subtle, hazy connections to indescribable emotions and ideas. When we abstract a concept, we sharpen it by removing some (if not most) of its properties, in particular these blurry edges. The resulting concept is said to be more abstract, and more general. (We might also call it more mathematical, though in very non-mathematical contexts this can sound as awkward as saying that soap is more delicious than urine.)

This abstraction and generalization serves at least two purposes. Firstly, a general concept has more **instantiations** than a specific one. The concept *three* , for instance, has only a numeric property, and hence is suitable for describing three dogs, three minutes, and three inches, so long as we only care about the number or amount of these things. And secondly, by eliminating and sharpening a concept's properties, our mind has less to worry about, and can focus more easily on just the properties that are relevant.

Forming abstractions from concepts is referred to as forming a **conceptual interface** . Abstractions are indeed interfaces: the resulting concepts are simpler and easier to work with, but because they are general they still apply to the hazy concepts we were originally interested in. This abstraction happens all the time without people knowing it. For example, suppose we are discussing the price of ice cream at a local store. Well, we have implicitly abstracted the original concept of their ice cream (which includes its price), by stripping away all properties except price. The taste of the ice cream, for example, is irrelevant to our current discussion. People don't recognize that they are making such abstractions, which is reflected in the fact that we often don't change the name of a concept when we abstract it for the sake of discussion: "They raised the price to $7.99 ? Sam's ice cream is getting out of control." "Well, at least it's not as bad as Ben & Jerry's !" . This can result in confusion called "mistaking the model for the thing being modelled" .

When the abstraction is so extreme that the new concepts are mathematical, this process is also known as **mathematical modelling** . This is a nice state of affairs, because it means that the whole power of mathematics is at your disposal for discussing the original concepts. (Unfortunately, it doesn't seem like there could be any useful mathematical model of love.)

Mathematics and precision

I've already mentioned how mathematical properties are more precise than nonmathematical ones. For instance, the human concept of *addition* has something to do with combining, putting things together in a certain way. This concept is abstracted into the mathematical concept $+$. But please note that $+$ is not completely precise, because even though we know many of its properties (for example, it is associative, symmetric, $1+1=2$, etc) , we don't know explicitly all the properties it has. For example, Goldbach's conjecture is a property involving addition (as well as other symbols), and nobody has any idea whether Goldbach's conjecture is true.

So to summarize, fuzzy human concepts are abstracted into fuzzy mathematical concepts. What has been gained? Well, quite a bit, because the properties of mathematical concepts are usually **completely precise** , or **formal** .

How can anything be completely precise, and have no ambiguity? Clearly if something is completely precise, it should have nothing to do with ideas, feelings, or anything semantic, because meaning is never completely precise. If we take away meaning, all that's left are raw symbols, or in other words, **syntax** . So a completely precise property would have to be formulated in terms of symbolic manipulation, like for instance, copying a string of symbols verbatim, or reversing their order, or by adding a symbol to the beginning of the string, etc etc.

And indeed, many mathematical properties are completely formal. For instance, $+$ has the following property:

$$\text{for all } x \text{ and } y, \quad x+y = y+x \quad .$$

This is known as the 'symmetry' of $+$. To my personal tastes, this formal property is an accurate translation of a fuzzy property of the fuzzy human concept of addition. In other words, $+$ is an appropriate abstraction or generalization of *addition* , at least as far as symmetry is involved.

But let's look at this property of symmetry: it tells us we can take the left argument of $+$ and the right argument of $+$ and switch them. This is a raw syntactic operation, a completely formal property. It is, in a sense, meaningless.

And even though it is meaningless, it is highly useful. If we are faced with a fuzzy human situation involving addition, we simply model the addition with $+$, and now we can confront the model of the situation in a completely precise way, just by moving symbols around the page. And at the end of the day, when we look to see what mathematics has told us about the original situation, our raw syntactic manipulations will have significant semantic consequences indeed.

Designing interfaces

So, to reiterate, a human, mathematical concept like $+$ is highly precise in that we can give its properties formally, in terms of symbol manipulation; but imprecise in that we don't know explicitly all its properties. In calculational mathematics, we often take abstraction even further: the concepts we deal with may have just one or two formal properties, and nothing else. As far as I can tell, this is as precise as a concept can be.

We usually design such concepts when attempting to solve a problem. First we translate the problem into a formal notation, often using traditional symbols like $+$, but sometimes using completely undescriptive ones like f . Then we forget completely about all the properties these concepts should have: we start from scratch. (Often, to help ourselves forget, we change the name of the symbol to something unfamiliar.)

Instead of starting from the properties we know, instead of starting from our previous knowledge, we focus on the problem itself, which is now given in terms of symbols. We use analysis and design to discover what properties of the symbols are needed to solve the problem, and then **postulate** those properties, declaring them by fiat. (This process was dubbed **the nabla trick** by Wim Feijen.)

Now, we have to play this game with a bit of care. For instance, let's say the original problem involved *subtraction* , which we modelled with the mathematical concept $-$. Then, we wiped the slate clean, and forgot all its properties. In fact, we really wanted to help ourselves forget, and so renamed $-$ into $\diamond$. So far so good. Then we began working on the problem, and discovered that we would like to have the following property of $\diamond$:

$$(*) \qquad \text{for all } x \text{ and } y \text{ ,} \qquad x \diamond y \;=\; y \diamond x \qquad .$$

With this property in hand, we were able to solve the problem. That is, we solved the problem in terms of the highly abstract $\diamond$, which has only one property.

Now we would like to use this result in order to solve the original problem, and to do that we need to **instantiate** $\diamond$ with the intended interpretation, the mathematical concept $-$. But there is an insurmountable difficulty, because $-$ does not have the property given in $(*)$. (For example, $2-3 \;\neq\; 3-2$.) This is a consequence of being very naive: when we were working on the very abstract version of the problem, we completely forgot about the intended application. We made an interface which was very easy to work with, but was not applicable to the original concepts.

One lesson to be learned from this is simply that we should **postulate** with care, because careless postulation of properties may yield a useless result. But another lesson to be learned is that we can very often discover new results in this way, that have nothing to do with the original problem. For example, although $-$ does not have property $(*)$, many many concepts do, like $+$. This means that the original problem has been solved, provided we replace *subtraction* by *addition* , or by any other concept with property $(*)$. In other words, we have a highly general, and hence useful, result… just maybe not useful in the way we wanted it to be.

The postulational method is not limited to mathematical contexts, either. If we wanted to discuss *love* , for example, we might aim to discover just what about *love* we need to know in order to carry out our discussion. This opportunistic approach is an excellent way to get a firm grasp on a hazy concept.

So to sum up, the postulational method and the use of highly highly abstract concepts are quite useful. They give us very general results, and help us discover just what properties are needed to solve a problem. And when we don't know any properties of the concepts/symbols, this approach can help direct our investigation into what properties they have.

Written: Culver City, 25 August 2006
Revision: Santa Cruz, 17 October 2006

Jeremy Weissmann
11260 Overland Ave. #21A
Culver City, CA 90230
USA
jeremy@mathmeth.com