# **Arguing About Radioisotope Dating**

# Laura Rassbach\*, Elizabeth Bradley\*, Kenneth Anderson\*, Christopher Zweck^, and Marek Zreda^

\* University of Colorado Department of Computer Science Boulder, CO 80309-0430 laura.rassbach@colorado.edu, lizb@cs.colorado.edu, kena@cs.colorado.edu

#### Abstract

We present a prototype of AICronus, an argumentation system that automates a challenging reasoning process used by experts in cosmogenic isotope dating. The architecture of the system is described and preliminary results are discussed.

# **1. Introduction**

Scientific reasoning is a complex process, alternately requiring flashes of insight and tedious analysis. This dichotomy is evident in the process of determining an age for a landform using cosmogenic isotope dating. Experts in the field of cosmogenic geochronology frequently spend months on repetitive mathematical tasks, until they have gathered enough information to suddenly understand the data. The AICronus project is aimed at understanding and automating this process.

Automating this reasoning process is challenging because the science of cosmogenic isotope dating is quite new. There is only a relatively small number of completed, detailed analyses to draw knowledge from. Therefore, it is necessary to build a knowledge base through interaction with experts. Unfortunately, it is often difficult for experts to clearly articulate how and why they come to specific conclusions. We have been working with experts in this field for more than two years, but our knowledge base is still very incomplete. Complicating the difficulty of acquiring the knowledge used by experts, the reactions that produce cosmogenic isotopes are not fully understood, leading to incompletely understood theories about them. As a result, most expert analysis relies on heuristics that are sometimes imprecise and frequently contradictory. Any system that automates the geological timeline construction process, then, must be able to handle both contradiction and uncertain heuristics gracefully.

AlCronus addresses these issues through the use of a nonmonotonic logic called argumentation. Argumentation uses symbolic logic, so that rules acquired from experts can be directly input into the system. In addition, the system's reasoning can be presented to the user in a legible <sup>^</sup> University of Arizona
 Department of Hydrology and Water Resources
 Tucson, AZ 85721
 czweck@hwr.arizona.edu, marek@hwr.arizona.edu

format, facilitating engagement and speeding further knowledge engineering. In the argumentation architecture used by AICronus, conclusions can receive partial support (modeling uncertain heuristics) and support for a conclusion can be defeated by contrasting evidence or rules (handling contradiction gracefully).

This paper presents the prototype version of the AICronus system along with some preliminary results, which show significant initial success in accurately modeling the reasoning process of isotope dating experts. Section 2 details the process of constructing a chronology for a landform using cosmogenic isotope dating. Section 3 discusses the particular challenges that arise in attempting to automate parts of this process. Section 4 demonstrates how argumentation addresses these challenges. Section 5 discusses the AICronus architecture in more detail. Section 6 walks through a concrete example of the working system. Section 7 discusses future work for the AICronus system, and section 8 covers related work.

# 2. Cosmogenic Isotope Dating

Cosmogenic isotope dating is a method for computing the age of a landform using cosmogenic isotope measurements of samples taken from that landform. Other methods for landform dating rely heavily on heuristic examination of features such as lichen growth (Bradwell 2001) or on dating materials associated with the lanform such as fossils (Noller, Sowers, and Lettis 2000). Cosmogenic isotope dating is more consistent and less subject to influence from the preconceptions of individual geologists.

This dating procedure is based on the knowledge that cosmic ray particles hit the earth at a fairly constant rate. When these particles come into contact with certain target elements, they can change them into cosmogenic isotopes (e.g. Chlorine-36 or Aluminum-26). The creation of these isotopes happens at a calculable rate. Most types of cosmic rays penetrate only a few meters, so the radioactive isotopes are generated almost exclusively at the surface. This knowledge enables a geologist to determine how long a particular sample has been at the surface based on the number of radioactive atoms present, the sample's chemistry, and other factors. The mathematics involved are



Figure 1: The process of assigning a landform's age.

quite complicated, and are handled by a different system being developed by the iCronus project at CU Boulder (Anderson and Bradley 2006).

For many landforms, the length of time a sample has been at the surface is actually a measure of the age of the landform (e.g. moraines, which are formed by glaciers carving boulders from deep underground and eventually depositing them, along with soil, as the glacier retreats). Other landforms are formed over a longer period of time, or from rock that was at the surface prior to the landform's formation. In these cases the length of time that samples have been at the surface can provide other information (e.g. how long the landform took to form) but will not give the actual age of the landform. Currently, cosmogenic isotope dating is used primarily to estimate ages of suddenly-created landforms.

The process of cosmogenic isotope dating begins with an expert taking samples of surface rock—generally thin chips from several boulders—from a single landform. Significant expertise is needed to choose good samples: sample boulders should not be excessively weathered, should show no signs of having been rolled or turned, should usually be of similar composition to the surrounding surface, etc. In many cases it will only be possible to take a small number of samples that meet these requirements. Experts also record as much data as possible about the location and status of the samples, including the sizes of the boulders they are collected from, the amount of visible sky, and the sample's exact location.

After collection, the expert sends the samples to an accelerator mass spectrometry (AMS) lab that measures the chemical properties of the samples. This includes the chemical composition and the percentage of certain isotopes (e.g. Chlorine-36 compared to overall Chlorine). The lab's services are extremely costly, further limiting the number of samples for which data are available. Based on these measurements, information about where the sample was taken, and a large amount of background knowledge, the expert calculates preliminary (or "apparent") ages for

the samples. The background knowledge involved includes data about changes in cosmic ray intensity, changes in sea level (which affects cosmic ray intensity at particular altitudes), and information about the production rates of the isotope in question from other isotopes (Gosse and Phillips 2001). For many of these background data, multiple measurements are available and will yield different results. Handling these calculations and organizing the background data is the task of the iCronus project (Anderson and Bradley 2006).

Next, the expert compares the preliminary individual ages. If all the apparent sample ages for a single landform are the same, within the margins of error introduced by the AMS analysis, that age is assigned to the landform and the process is complete (the expected distribution is a Gaussian, the typical expectation for events affected by random errors). However, this happy situation rarely occurs. It is more usual for preliminary age measurements for different samples to differ by as much as 10,000 years (Shanahan and Zreda 2000). In this case, the expert attempts to explain the divergence so that s/he can assign a single age to the landform. The AICronus system is designed to assist with this explanation.

Sometimes there is no good explanation, or there are several explanations that cover the data equally well. Often the result of this first round of analysis is the conclusion that more samples are needed. In this situation, experts often reason from the existing data to guide further sampling–for instance, when it is not possible to determine whether erosion or inheritance caused a spread in ages, experts may take a soil sample or sample a landform at depth to obtain more information. This reasoning is not captured by the prototype version of AICronus described here; ideas for adding it to the deployed version are discussed in Section 7. Figure 1 illustrates this cyclic process.

# **3.** Automation Challenges

Most explanations for spread in apparent ages come from a short list of about fifteen geologic processes that affect the preliminary exposure times of samples from a single landform. For example, erosion gradually exposes new surfaces, causing some samples to have apparent ages much younger than the age of the landform. A "process" called inheritance affects samples that were exposed before the landform in question was formed, giving them apparent ages older than the age of the landform. Other processes include cover by snow or vegetation, gradual formation such as from soil deposits, or earthquakes, which may suddenly expose large amounts of rock at the surface. Multiple processes may act on a single landform. Finally, possibilities like lab error and mis-sampling must also be taken into account.

Although we need only consider explanations from among a relatively small number of processes, the complexities of how these processes affect the data especially when multiple processes are involved—make this task far from simple. Data are noisy and frequently



Figure 2: Apparent-age distributions (number of samples vs. apparent age) produced by various processes.

cannot be trusted (experts may have mis-identified the type of landform they are considering, for example), and the manifestation of one process may be quite similar to the manifestations of other processes.

Many processes that affect apparent ages of samples give the apparent age distribution a characteristic shape. When no geologic process has affected the samples, random errors in measurement are expected to produce a Gaussian distribution, with the true age of the landform equal to the peak of the curve. Matrix erosion of a moraine, which gradually exposes new boulders as the top soil of the moraine erodes, looks something like a skewed bell curve with the peak towards the older end of the scale. On the other hand, inheritance usually involves a simple uniform distribution over an age range. Figure 2 shows some examples of these distributions. Unfortunately, it is difficult to diagnose the process that affected a particular landform from the distribution of apparent ages because we rarely have enough samples to see the distribution shape. Instead, experts usually perform this diagnosis using heuristics about how various types of landforms are formed and how each process affects different landforms. Other heuristics are applied to the apparent distribution of the small number of samples: experts label three samples that are approximately evenly spaced erosion-not inheritance, as we would expect based on the a priori knowledge of sample distributions. The experts we are working with have not yet been able to explain this apparent contradiction.

Choosing which geologic process is responsible for the data is complicated by the fact that there is generally some evidence both for and against several processes, a major challenge for automatic analysis. From our observation of experts in this field, it appears that a standard approach to this problem is to select one process and look for evidence both for and against that process. If it is possible to gather enough evidence in favor of a process and not possible to gather a similar or greater amount of evidence against the process, then it is considered a good candidate for explaining the data.

# 4. Solution: Argumentation

The specific task of AICronus is to assist with the analysis of apparent sample ages and help determine what processes are good candidates to explain the spread in apparent ages. As just mentioned, many of the heuristics that experts use in this process are vague and sometimes contradictory. For example, matrix erosion is expected to produce a skewed bell curve of initial sample ages. In practice, however, experts usually assign matrix erosion to cases that have a uniform distribution (Shanahan and Zreda 2000). To further complicate matters, inheritance is the process expected to produce a uniform distribution! Contradictions may also arise when input observations are incorrect in some way, for example when samples are entered as members of one landform but have actually come from two different landforms (Desilets and Zreda 2006). Heuristics like "this is a moraine, so inheritance is more likely" are also common and are clearly not absolute; we do not always conclude inheritance when the landform is a moraine. Therefore AICronus must gracefully handle both contradiction and partial support.

In addition to these technical issues, experts are unlikely to agree with any conclusions made by AICronus unless they understand the reasoning behind those conclusions. Thus it is also critical to the usefulness of AICronus that it be capable of convincingly presenting the reasons for its conclusions. This capability provides the additional benefit that students of geology can examine the reasoning and heuristics that are used in selecting a process.

Argumentation systems are a good solution here. They provide the functionality needed for AICronus to be useful to both experts and students in cosmogenic isotope dating. They are capable of handling contradictory rules and input data, partial support for conclusions, and can report their reasoning in a clear and understandable way (Krause, Ambler, Elvang-Gøransson, and Fox 1995) (Doyle 1983). In fact, the reasoning used in argumentation appears to closely match the flexibility and methodology that experts in the field actually use in their analyses.

The argumentation framework used by AICronus is based on the Logic of Argumentation introduced by Krause et. al. (Krause, Ambler, Elvang-Gøransson, and Fox 1995). Unlike in traditional first-order logic systems, rules, input data, and "proofs" in argumentation systems may all be considered defeasible. Proofs in classical logic correspond to arguments in these systems—an argument is a reason for believing some conclusion, but contradictory arguments may also be formed.

Krause et. al. implement an argumentation system as an extension to a limited form of first order logic, where rules and data are labeled with a confidence level used to determine which of two arguments is stronger. As arguments are built, the confidences propagate to their conclusions using some system of combination (defined as appropriate to individual problems). Arguments are labled with the rules and data used to form them. The logic is limited because the 'not' operator is not implemented and disjunction is handled only implicitly by defining multiple rules that can apply to the same conclusion.

The confidence values in AICronus are in the range [-1, 1], and are currently combined using several different functions, selected by which rule is being used. Negative confidence in some literal is interpreted as confidence in the negation of that literal (zero confidence implies the system knows nothing about a term). Conclusions generated by the system are labeled with the arguments for and against them, so that as new information is discovered, the arguments about a conclusion can be examined and possibly defeated. An argument can be defeated in two ways: a stronger argument can be found against the conclusion of the argument (rebuttal), or arguments can be found against the evidence used in the defeated argument (undercutting).

AlCronus treats the arguments for and against a particular conclusion like grains of sand on a scale. Stronger arguments, formed using rules and data with higher confidence levels, add more weight to their side of the balance. However, a large enough number of poor arguments on one side of the scale can overpower a single good argument on the other side. Unlike a balance loaded with sand, additional "weight" is added in a system of decreasing returns: two poor arguments of the same quality have less weight than one argument of twice their quality. That is, if a single argument has a confidence of 0.8, it will defeat two combined, rebutting arguments, each with a confidence of 0.4 (but it will be defeated by three such arguments). Undercutting is handled by reducing the degree of confidence in the undercut argument.

# **5.** Constructing Arguments

AICronus takes as input all available data about a set of samples, along with information about the site where they were collected. This includes both qualitative data (e.g. the type of landform and the color of the boulder the sample was taken from) and quantitative data (e.g. the calculated apparent sample ages and the elevation of the landform).



Figure 3: An example AICronus argument. At the top is the conclusion being argued about. Beneath is a collection of trees arguing about this conclusion. Rules are shown in boxes and entered observations in ovals.

Information about nearby landforms may also be included in the input, since the ages of these landforms may imply strict upper or lower bounds on the age of the landform in question (e.g. moraines must decrease in age as the elevation in a single valley increases because of the way they are formed).

AlCronus generates a list of processes that may have affected the landform, with more-common processes (as specified by experts) higher on the list. Arguments for and against each process on the list are generated via backward chaining, building an argument tree. Once a process has been found for which the "pro" evidence significantly outweighs the "con" evidence, the system stops and reports its results to the user. These results include all of the processes so far considered and the complete arguments for and against each process. Processes with the most convincing arguments are listed first. The user can choose to generate arguments for more processes if s/he finds the presented results insufficiently convincing.

#### 5.1. Arguments

Rules in AICronus have a standard first-order logic structure, where a rule is written in the form A => C. A may be a single literal or the conjunction of several literals. An argument is a collection of trees, with rules from the system's database forming the nodes of the trees. Rules in child nodes have the same variable in their conclusions as one of the literals on the antecedent side of the parent node's implication. At the root of each tree is a single rule that allows AICronus to argue about whether some particular process is responsible for the observed data. The leaves of the trees are drawn from the observations entered by the user.

Figure 3 shows an example collection of argument trees in AICronus which might read: "Erosion is a likely explanation because moraines are likely to erode and this landform is a moraine. However, there is no visual evidence of erosion such as a flat crest or weathering, making erosion a less convincing conclusion." The total confidence in the argument is determined by the total confidence in the trees that argue for the root process versus the total confidence in the trees that argue against the root process.

#### 5.2. Rules

AICrounus rules have a number of parts. Rules are implicitly built around the classic implication structure from first-order logic. In addition, they contain guards to indicate when a rule is applicable, and instructions for how to obtain a confidence for the conclusion of the rule from the confidences in its premises. Figure 4 shows an example of an AICronus rule in with its parts labeled. Each of these parts is discussed in more detail below.

**Conclusions and Antecedents** These two fields define the implication that is the main part of the rule. In a classical first-order logic system, when the antecedents of an implication are true, we can conclude the conclusion

```
(
    rule conclusion
    matrix-erosion
    guard
    (landform-type = moraine)
    antecedents (or premises)
    ((calculation
        erosion-choice := get-maximum age)
    (argument
        evidence types
        good-final-choice erosion-choice))
    confidence combination
    (selector 2 .8)
}
```

# Figure 4: An example AICronus rule with individual parts labeled.

with absolute confidence. In AICronus, when there are arguments supporting the premises in a rule, they can be used, along with the rule, to form an argument for the rule's conclusion. Unlike in a classical system, this argument may eventually be overturned, possibly causing us to conclude the negation of the conclusion.

When combining antecedents with a rule to form an argument, the backwards-chaining engine makes no distinction between evidence for the premises and evidence against them. Antecedents with a negative confidence rating are treated identically to those with a positive confidence, although frequently (as in the example) this case will generate an argument against the rule's conclusion.

The rule in Figure 4 encodes the fact that when matrix erosion is the selected explanation for a spread in sample ages, we will choose the oldest sample age as the correct age for the entire landform. It is important to make sure that this age will be compatible with any restrictions on this landform's age due to known ages of surrounding landforms. In addition, our experts use a heuristic that tends to reject the choice of a single sample's age for the age of the landform when it is distant from the other samples. Arguing for whether the chosen age is a "good-final-choice" addresses both of these issues. Each of the antecedants is labeled with a type to tell the engine how to handle it. These types are discussed in more detail in section 5.3.

**Guards** The guards on an AICronus rule prevent the system from building arguments using rules that are not applicable to the current case. For example, AICronus has a rule that snow cover is more likely if samples appear younger at higher elevations. However, elevations are recorded for all samples, even when they were collected at essentially the same elevation. Obviously the rule only makes sense when we are dealing with elevation ranges large enough to have different levels of snow cover. Therefore the guard on the rule states that it is only applicable when the elevations of the samples have a large enough range.

In Figure 4, the guard states that the rule is only applicable when the landform under consideration is a

moraine. This is because matrix erosion is a process specific to moraines, and it would not make sense for the system to consider it as a possible conclusion on a landform that is not a moraine.

Confidence Combinations We intend to produce a standardized methodology to combine argument consistently meaningful confidences into values. Unfortunately, we currently have an insufficient number of cases to generalize confidence. To allow for rapid feedback and prototyping, confidence combinations are handled somewhat individually until we can determine the correct unified method. We add new combination methods to the system as we add rules for which none of the current methods seem adequate, so the list of combinations types is expanding. Our current methods for confidence combination include:

- Scalar combinations: this method uses a linear combination to combine the confidences in the antecedents into a confidence in the rule's implicant. These combinations are used in rules where all the antecedents are directly related to the conclusion. For example, moraines are formed with a pointed crest which flattens as they erode: A flattened moraine crest is evidence for matrix erosion, and an unflattened crest is evidence against it. Selectors (such as in Figure 4) are a special case of scalar combinations where all weights but one are 0.
- Asymmetric scalars: this combination is like simple scalars, except the linear combination coefficients change based on whether the confidence in the antecedent is positive or negative. These are used in cases where the antecedent is more useful in drawing conclusions one way than another. For instance, we may be interested in whether one sample came from a different landform than the rest. If the samples were collected from the bedrock of the area, we can be very certain that they came from the same landform. However, we cannot be confident of a different origin simply because the samples were not taken from bedrock.
- And-like combinations: if the confidence of every antecedent is positive, a constant confidence is assigned to the implicant of this rule. If any antecedent's confidence is negative, then the confidence is the negative of the constant. For example, if all samples entered have similar ages, we can conclude no process is needed to explain the data. Otherwise, we need to look for some process to explain our observations.
- Combination combinations: some rules use a composition of the other combinations (e.g. a scalar combination, instead of a constant, as the confidence value for an and-like combination). For instance, we can guess that a location is not cold enough for very much snow, at least in recent geological time, if it is both near the equator and at a relatively low elevation. Our confidence in the likelihood that the area does not get cold enough

for significant snowfall goes up as we move closer to the equator and to even lower elevations.

It is apparent from this list of combination mechanisms that most combinations include some scaling factors. Selecting the correct weights for these scaling factors is problematic from both a theoretical and practical standpoint (Doyle 1983). We have addressed this problem initially by using only a small set of actual values for weights corresponding to English phrases used by experts to describe how convincing they find the evidence in each rule (such as "very convincing," "somewhat convincing," "minor evidence," etc.). We have assigned each of these phrases a weight value. Numeric weights are used in the AICronus prototype solely for convenience purposes, and may be changed to symbolic confidence values when we begin to use a more-uniform system of confidence combination. We believe one of the strengths of argumentation is that it allows the system-and its usersto focus on the reasoning behind a conclusion rather than the specific value assigned by the system, so that the somewhat *ad hoc* choice of values is not a major weakness.

#### 5.3. Evidence

The data AICronus uses to draw its conclusions are referred to as evidence. The antecedents in a rule's implication are patterns for evidence—they indicate what evidence will be needed to satisfy the rule. The actual data that causes AICronus to conclude something about the antecedent is the evidence. The system has four different kinds of evidence: observations, simple calculations, complex calculations, and arguments. The distinction between these types of evidence is inspired by the PRET (Stolle and Bradley 1996) system. The separation allows less computationally intensive rules to be considered first.

**Observations** Observations are direct uses of the data entered by the user. Usually an observation is some binary involving the data, for example checking that all samples have apparent ages less than a certain value. Because the user's observational evidence has more confidence if the relation is stronger. For an antecedent like "elevation < 10000 ft.", we will be more confident that the condition has been met with an elevation value of 7000 ft. than a value of 9999 ft. Observations may also take the form of a quantifier such as for-all or there-exists. These are handled by selecting the highest (for there-exists) or lowest (for for-all) individual confidence value among the quantified entities.

**Simple Calculations** Simple calculations are generally calculations of simple statistical properties of entered data. They are used for the purpose of generating the calculation's results and all simple calculations have a confidence value of 1. A simple calculation might find the mean of all apparent sample ages so that another part of the rule can check that all apparent sample ages fall within a certain distance of this mean.

**Simulations** More-complex calculations are called "simulations" because they usually are. Simulations have varying confidence values based on their results. They are implemented as procedures called by the engine examining the rules, allowing them to be as complex as necessary. An example simulation tries different levels of erosion, looking for the rate that best explains the distribution of the apparent sample ages. The simulation returns this erosion rate (which can then be checked to confirm, e.g., that it is reasonable for the climate of the sampling area) and a confidence value indicating how well the returned rate reduces the spread in the calculated ages.

**Arguments** Sometimes the antecedents of a rule cannot be directly gleaned from the input data. In this case it may be necessary to build a sub-argument for an antecedent and to use the sub-argument as evidence. For example, we know that snow cover is much less likely in areas that are not cold. The system can build a sub-argument for whether the sampling area is cold as part of an overall argument about snow cover.

# 6. AICronus in Action

Although still in a prototype stage, AICronus is able to produce answers and arguments similar to those produced by experts. Here is an example set of input data created for the purpose of communicating with experts about their reasoning about specific examples.

Landform Type: Moraine Flat Crest

Sample	Age (yrs)	Error (yrs)	Chemistry
1	95000	5000	А
2	100000	6000	А
3	105000	4000	А
4	110000	4500	A
5	115000	5500	A

Two experts, shown this set of input data, concluded that the process affecting the data was almost certainly matrix erosion, primarily because of the distribution of apparent ages but also because the landform is a relatively old moraine with a flat crest. AICronus considered inheritance, but rejected it because the errors were too small and because all the sample appeared to have the same origin, making different inheritance levels for different samples unlikely.

Here is AICronus's actual output, given this input set:

```
argument for conclusion matrix erosion:
total confidence: 0.87
evidence for erosion:
    age is approximately linear
    landform is relatively old (>100000 yrs)
    visual erosion observed
    argument for conclusion visual-erosion:
```

```
total confidence: 0.6
        evidence for visual-erosion:
          flat crest
        evidence against visual-erosion:
           (none)
     consistent with other landforms in area
      argument for conclusion consistent-age:
      total confidence: 1
       evidence for consistent-age:
         no other landforms known
       evidence against consistent-age:
          (none)
     landform is a moraine
   evidence against erosion:
       (none)
 argument for conclusion inheritance:
 total confidence: 0.37
   evidence for inheritance:
     consistent with other landforms in area
      argument for conclusion consistent-age:
      total confidence: 1
      evidence for consistent-age:
        no other landforms known
      evidence against consistent-age:
         (none)
     removal of older samples allows "no
process"
     landform is a moraine
     does not violate max theoretical inher-
itance
   evidence against inheritance:
     small error values (<10000 yrs)</pre>
     samples have same origin
      argument for conclusion same-origin
      total confidence: 0.5
      evidence for same-origin:
        all samples have same chemistry
        not taken from bedrock
      evidence against same-origin:
        landform is a moraine
```

Because AICronus rules are already built around the idea of arguments, producing an English argument from a set of rules is relatively simple. The system simply needs to print a representation of all the argument trees for all the conclusions it considered. Since the trees already contain both the rules and the evidence that allowed the system to use each rule, simply translating each rule and its evidence into English produces a complete and coherent argument.

Although both inheritance and erosion have positive confidence values, the system's confidence in erosion is much higher. This exactly matches the judgement of the experts who were shown these data. Moreover, AICronus's arguments about the possible processes closely match the arguments given by the experts in each case. Despite the difficulties inherent in the field of cosmogenic isotope dating, AICronus already shows significant promise in understanding and automating the reasoning used by experts.

# 7. Future Work

AICronus is a work in progress. We plan a number of improvements over the next several years. The most critical of these improvements is expanding the system's knowledge base. We are in the process of using this prototype version to solicit feedback and new knowledge from experts. In addition, we are working on integrating this system with the iCronus project (Anderson and Bradley 2006) so input data need not be entered by hand and output arguments can be presented visually rather than via the current command-line interface. Our current interface requires the user to enter a complete set of samples as Scheme code—an obvious drawback for many users. We expect the system to go into regular use by geologists once these improvements are complete.

We are considering other improvements to make the system more user-friendly. These include removing the numeric confidence values in the output and presenting arguments in a more natural prose form. Currently the system does not provide any assistance for going back to collect more samples to distinguish between processes that appear to have equally good arguments. We are considering an approach similar to (McIlraith and Reiter 1992) for implementing this functionality.

Other future projects include allowing the user to engage in an argument with the system to update the knowledge base over time based on user input. We also hope to apply this framework to problems in other fields. As discussed above, the major strength of the argumentation framework is that it copes well with uncertain and contradictory information. These features often appear in new or rapidly changing scientific fields. Although we expect that the rule and argument structures of AICronus will translate well between fields, areas with more conclusions might require a more performance-aware engine. Other fields might also require us to consider new types of evidence.

# 8. Related Work

Many diagnostic systems solve problems similar to the one solved by AICronus, in which there is some normal, expected behavior (in isotope dating, all samples of the same apparent age) and the causes of divergences from this behavior (e.g. a geologic process) must be diagnosed. However, the predominant paradigm in medical diagnosis is to build a complete model of a system and to use that model to make predictions about malfunctions (Lucas 1997), (Struss 2004). This methodology is not suited to our particular domain because complete models of most geologic processes simply do not exist. In addition, modelbased systems are not as suited to handling contradiction.

Diagnosis systems that handle contradiction do exist, for example (Doyle 1983), (Santos 1991), (Cem Say 1999) and

(Gaines 1996). However, all of these systems use "absolute" rules. It is not possible to express the idea that some data may only partially support a conclusion. Instead, the conditions under which the rule does not provide support are explicitly encoded in the system. AICronus needs to include rules for partial support of conclusions in order to accurately reflect the reasoning process used by experts. For instance, experts are more likely to accept (and require more evidence to reject) an "inheritance" conclusion for a moraine than for other landform types. Trying to model this behavior without an ability to express partial support would be extremely difficult. On the other hand, all of these systems are capable of presenting their reasoning to the user to help convince experts of initially rejected conclusions, an important feature of AICronus. (Santos 1991) is also capable of presenting alternative conclusions to the user so that if the user does not agree with a particular conclusion the tool is likely to still be useful (another AICronus feature).

Several authors have discussed the virtues of presenting the reasoning behind a system's conclusions in the form of trees or arguments, including (Boy and Gruber 1990) (Bouwer and Bredeweg 2002) and (Gaines 1996). (Puyol-Gruart, Godo, and Sierra 1992) points out that even when a particular conclusion cannot be reached by a reasoning system, it is likely that presenting what the system *has* managed to determine will be useful to the user. AICronus handles this situation by reporting its complete arguments, even in cases where the absolute values of the confidences are quite small.

Case-based reasoning (Kolodner 1993), (Cunningham, Bonzano, and Smyth 1995), and (Clark 1989) presents a way to sidestep the issues of partial support and contradiction by presenting intact the reasoning of experts on previous cases that are similar to the current problem instance. Unfortunately, case-based reasoning is unsuitable to AICronus because the field of cosmogenic isotope dating is still very new and relatively small. As a result, there are too few already-analyzed cases to cover all of the possible variables in selecting a responsible process. (Surma and Vanhoof 1995) offers a solution to this objection by using rules for "normal" cases and case-based reasoning for cases that are exceptional in some way. Unfortunately, the problem being solved by AICronus has so many variables to address that it is difficult to classify any case as "normal."

(Turner 1992) uses schemas—abstracted cases—to perform diagnosis by considering particular symptoms. When a symptom is unique to a particular type of disease, the system considers diagnosing that disease. If the symptoms expected for that disease are observed, then it is considered a correct diagnosis. This architecture is not suitable for isotope dating because it fails to handle contradiction well. In addition, schemas are difficult to extract because it is difficult to determine what is typical for any process.

Several kinds of defeasible reasoning besides argumentation have been put forth by various authors. These include circumscription (McCarthy 1980), (McCarthy 1986), default reasoning (Reiter 1980), (Doyle 1983), and other forms of nonmonotonic reasoning (Pereira, Alferes, and Apar'icio 1991), (Gaines 1996). Circumscription allows the definition of normal situations and the cases that can circumscribe them. It requires the definition of specific aspects that are abnormal only in abnormal situations, so that it is necessary to create a large number of "aspect" variables to express all of the possible abnormal situations. Default reasoning uses rules with default conclusions and then defines specific exceptions where they do not apply. This is similar to the "guards" on AICronus rules which prevent them being used to build arguments in some situations. The nonmonotonic logic defined in (Pereira, Alferes, and Apar'icio 1991) assigns a likelihood to various rules so that they can normally, sometimes, or exceptionally apply. Rules have conditions stating specifically when they do apply. (Gaines 1996) uses a tree structure for rules with default conclusions at the root and repeated refinements or rejections of the initial conclusion(s) as the tree is descended.

While all of these logics are excellent choices for solving many different problems, they all require some explicit definition of when particular rules are defeated. The heuristics used by our experts are insufficiently complete for these explicit definitions. Also, all of these nonmonotonic logics use defeat of specific rules rather than attacking conclusions. AICronus rules are not bound in a strict fashion to conclusions; a rule may be in support of a conclusion (but turn out to be unimportant in light of other rules or conclusions) or against one (but be negated by the presence of higher-confidence results elsewhere). (Etherington, Kraus, and Perlis 1991) describes other problems with various nonmonotonic logics.

There is a large body of work on different kinds of argumentation systems. Most of this work grapples with the question of when it is appropriate to declare a particular argument defeated, with different authors reaching various conclusions. Most authors (Dung 1995), (Pollock 1994), (Vreeswijk 1991), (Farley 1997), and (Prakken 1996) consider only absolute defeat of arguments. Little work on partial support and defeat has been done, although the Logic of Argumentation introduced by Krause et. al. (Krause, Ambler, Elvang-Gøransson, and Fox 1995), on which the AICronus framework is based, does partially address these issues.

Few results exist for applying argumentation to specific problems. Most practical systems are aimed at communication-based applications, especially communication between agents (Parsons, Sierra, and Jennings 1998). The idea of argumentation as a form of communication has also been explored by (Prakken 1996), (Farley 1997) and (Vreeswijk 1993), who cast the construction of arguments as a form of dialectics. In these systems, two agents repeatedly try to form arguments for a given conclusion, and then defeat those arguments. (Prakken 1996) allows defeat to take the form of defeating particular rules, rather than only the more-traditional undercutting and rebuttal. This defeat is analogous to the attachment of confidence values to specific AICronus rules; rules with greater confidence can defeat rules with smaller absolute confidences. (Farley 1997) allows the user to globally alter the relative strength of arguments. Three modes are allowed, where a conclusion is made if some argument for it exists, a conclusion is accepted if there are more arguments for it than against it, and a strict mode where a conclusion is believed only if there is an argument for it and all arguments against it are defeated. The second mode in particular is similar to the mechanism used by AICronus, except that the strengths of the arguments in (Farley 1997) are not determined by which rules are used to form them all defeasible rules have the same believability.

# 9. Conclusion

Although still in its prototype stage, AICronus is a promising model for the process of cosmogenic isotope dating. Using a logic of argumentation, we have generated preliminary results that closely parallel the reasoning and explanations of experts in the field. We expect that once the knowledge base for the system is complete AICronus will be able to reach insightful conclusions more quickly and consistently than experts under certain circumstances. In particular we expect this benefit in cases where superficially contradictory evidence disguises an extremely typical manifestation of some process.

We expect that AICronus will be a significant advancement for the field of cosmogenic isotope dating. Among other things, creating AICronus forces experts to make explicit many implicit rules and assumptions, allowing the easier identification of faulty or missing assumptions. These assumptions can then be used consistently between different experts. New knowledge about cosmogenic isotope dating can be written in an unambiguous form, allowing easier communication of new knowledge. Finally, we hope that AICronus will increase the speed of discovery of new knowledge about the ancient Earth.

#### **10.** Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. ATM-0325812.

#### References

Anderson, K. and Bradley, E. 2006. iCronus information is available at: http://serl.cs.colorado.edu/projects/icronus/

Bouwer, A. and Bredeweg, B. 2002. Aggregation of Qualitative Simulations for Explanation. In Proceedings of the Sixteenth International Workshop on Qualitative Reasoning. Barcelona. Boy G. and Gruber, T.R. 1990. Intelligent Assistant Systems: Support for Integrated Human-Machine Systems. In Proceedings of 1990 AAAI Spring Symposium on Knowledge Based Human-Computer Communication. Stanford University.

Bradwell, T. 2001. A New Lichenometric Dating Curve For Southeast Iceland. *Geografiska Annaler: Series A*, *Physical Geography* 83(3): 91-101.

Cem Say, A.C. 1999. Using Inter-Behavior Contradictions for Modeling and Diagnosis. In Proceedings of the Thirteenth International Workshop on Qualitative Reasoning. Loch Awe, Scotland.

Clark, P.E. 1989. Exemplar-based reasoning in geological prospect appraisal, Technical Report, 89-034, Turing Institute, Univ. of Stratchelyde.

Cunningham, P., Bonzano, A., and Smyth, B. 1995. An incremental case retrieval mechanism for diagnosis. Technical Report TCD-CS-95-01, Department of Computer Science, Trinity College.

Desilets, D. and Zreda, M. 2006. Elevation dependence of cosmogenic 36Cl production in Hawaiian lava flows. *Earth and Planetary Science Letters* 246: 277-287.

Doyle, J. 1983. Methodological Simplicity in Expert System Construction: The Case of Judgments and Reasoned Assumptions. *AI Magazine* 4(2): 39-43.

Dung, P.M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n-person games. *Artificial Intelligence* 77(2): 321-357.

Etherington, D., Kraus, S., and Perlis, D. 1991. Nonmonotonicity and the Scope of Reasoning. *Artificial Intelligence* 52(3): 221-261.

Farley, A.M. 1997. Qualitative Argumentation. In Proceedings of the Eleventh International Workshop on Qualitative Reasoning. Cortona, Italy.

Gaines, B.R. 1996. Transforming Rules and Trees into Comprehensible Knowledge Structures. In *Advances in Knowledge Discovery and Data Mining*, 205-226. Cambridge, MA: MIT Press.

Gosse, J. C. and Phillips, F. M. 2001. Terrestrial cosmogenic nuclides: theory and application. *Quaternary Science Reviews* 20(14): 1475-1560.

Kolodner, J. 1993. *Case-Based Reasoning*. San Mateo, CA: Morgan Kaufmann.

Krause, P., Ambler, S., Elvang-Gøransson, M., and Fox, J. 1995. A logic of argumentation for reasoning under uncertainty. *Computational Intelligence* 11: 113-131.

Lal, D. 1958. Investigation of Nuclear Interactions Produced by Cosmic Rays. PhD diss., Univ. Bombay.

Lucas, P.J.F. 1997. Symbolic diagnosis and its formalisation. *The Knowledge Engineering Review* 12(2): 109-146.

McCarthy, J. 1980. Circumscription—a form of nonmonotonic reasoning. *Artificial Intelligence* 13(1,2): 27-39, 171-172.

McCarthy, J. 1986. Applications of circumscription to formalizing common sense knowledge. *Artificial Intelligence* 26(3): 89-116.

McIlraith, S. and Reiter, R. 1992. On tests for hypothetical reasoning. In *Readings in Model-Based Diagnosis*, 89-96. San Mateo, CA: Morgan Kaufman.

Noller, J. S., Sowers, J. M. and Lettis, W.R., Eds. 2000. Quaternary Geochronology: Methods and Applications. AGU Reference Shelf. Washington, American Geophysical Union.

Parsons, S., Sierra, C., and Jennings, N. 1998. Multicontext Argumentative Agents. In Proceedings of the Fourth International Symposium on Logical Formalizations of Commonsense Reasoning. London.

Pereira, L.M., Alferes, J.J., and Apar'icio, J.N. 1991. Nonmonotonic reasoning with well founded semantics. In Proceedings of the Eighth International Logic Programming Conference, 475-489. Paris: MIT Press.

Pollock, J.L., 1994. Justification and defeat. *Artificial Intelligence* 67(2): 377-407.

Prakken, H. 1996. Dialectical proof theory for defeasible argumentation with defeasible priorities. In Proceedings of the Biannual International Conference on Formal and Applied Practical Reasoning Workshop, 202-215. Bonn.

Puyol-Gruart, J., Godo, L., and Sierra, C. 1992. A Specialization Calculus to Improve Expert Systems Communication. In Proceedings of the European Conference on Artificial Intelligence, 144-148.

Reiter, R. 1980. A logic for default reasoning. *Artificial Intelligence* 13: 81-132.

Santos, E. 1991. On the Generation of Alternative Explanations with Implications for Belief Revision. In Proceedings of the Seventh Conference on Uncertainty in Artificial Intelligence, 339-347. San Mateo, CA: Morgan Kaufmann.

Shanahan, T.M. and Zreda, M. 2000. Chronology of Quaternary glaciations on Mount Kenya and Kilimanjaro. *Earth and Planetary Science Letters* 177: 23-42.

Stolle, R. and Bradley, E. 1996. A Customized Logic Paradigm for Reasoning about Models. In Proceedings of the Tenth International Workshop on Qualitative Reasoning. Stanford Sierra Camp, CA.

Struss, P. 2004. Deviation Models Revisited. In Proceedings of the Eighteenth International Workshop on Qualitative Reasoning. Evanston, IL.

Surma, J., and Vanhoof, K. 1995. Integrating rules and cases for the classification task. In Proceedings of the First International Conference on Case-Based Reasoning, 325-334. Berlin: Springer Verlag.

Turner, R.M. 1992. A view of diagnostic reasoning as a memory-directed task. In Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society. Bloomington, IN: Cognitive Science Society.

Vreeswijk, G. 1991. The feasibility of defeat in defeasible reasoning. In Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning, 478-483. Morgan Kaufmann.

Vreeswijk, G. 1993. Defeasible dialectics: a controversyoriented approach towards defeasible argumentation. *Journal of Logic and Computation* 3: 317-334.