

Automatic Invention of Integer Sequences

Simon Colton, Alan Bundy

Mathematical Reasoning Group
Division of Informatics
University of Edinburgh
80 South Bridge
Edinburgh EH1 1HN
United Kingdom
simonco,bundy@dai.ed.ac.uk

Toby Walsh

Department of Computer Science
University of York
Heslington
York YO10 5DD
United Kingdom
tw@cs.york.ac.uk

Abstract

We report on the application of the HR program (Colton, Bundy, & Walsh 1999) to the problem of automatically inventing integer sequences. Seventeen sequences invented by HR are interesting enough to have been accepted into the Encyclopedia of Integer Sequences (Sloane 2000) and all were supplied with interesting conjectures about their nature, also discovered by HR. By extending HR, we have enabled it to perform a two stage process of invention and investigation. This involves generating both the definition and terms of a new sequence, relating it to sequences already in the Encyclopedia and pruning the output to help identify the most surprising and interesting results.

Introduction

An integer sequence is an ordered set of integers such as the square numbers: 1, 4, 9, 16, ... Integer sequences arise in many areas of mathematics, and comprise an important subject area. The Encyclopedia of Integer Sequences (Sloane 2000) is an on-line repository of around 54,000 sequences collected over 35 years by Neil Sloane, with contributions from many mathematicians. To allow a sequence into the Encyclopedia, Sloane stipulates that it must be an infinite sequence of positive integers which is well defined and interesting. This rules out any randomly generated sequences which have no formula and many dull sequences which have no interesting features. Each sequence is given an 'A'-number which uniquely identifies it within the Encyclopedia, for instance the square numbers have number A000290.

We have used the HR program (Colton, Bundy, & Walsh 1999) to invent new integer sequences worthy of the Encyclopedia. HR performs theory formation in domains such as number theory, graph theory and group theory by inventing concepts and making and settling conjectures. We present here extensions to HR's abilities in number theory. Firstly, number theory concepts are presented as integer sequences, eg. the concept of an integer being prime is converted into the sequence of prime numbers. Next, taking advantage of the natural ordering of integers, we have given HR new ways to produce number theory concepts. Finally, we have enabled HR to provide justification why a new sequence

is worthy of the Encyclopedia by relating it to sequences already in the Encyclopedia. Sometimes the relationships found are surprising and non-trivial to prove which adds to the interestingness of the new sequence. HR therefore employs a two step process to find new sequences:

- A sequence is invented by generating a definition and determining the first terms of the sequence.
- The sequence is investigated by relating it to ones already appearing in the Encyclopedia.

Invention of Sequences

The HR Program

HR is named after mathematicians Hardy (1877-1947) and Ramanujan (1887-1920). It is designed to model how mathematical theories can be formed from only the most fundamental concepts of a domain, such as addition and multiplication in number theory. HR is supplied with (a) some objects of interest such as groups, graphs or integers (b) some ways to decompose these into sub-objects, such as graphs into nodes, and (c) some relations between the sub-objects, such as nodes being adjacent.

Each initial concept is supplied with a data table of rows which satisfy a predicate, and a definition for the predicate. The first column of every data table contains objects of interest, the other columns contain sub-objects or integers calculated using the objects of interest and their sub-objects. For example, the concept of divisors of integers is supplied with the first data table in figure 1, where each row is an integer and a divisor. The definition supplied is: $[n, a] : a|n$, read "a divides n", which is used when generating a definition for any concept based on divisors.

Given the initial concepts, HR uses general production rules to turn one (or two) old concepts into a new one. For instance, the 'forall' production rule finds objects where a particular relation between its sub-objects is true in every case, eg. if the relation is adjacency of nodes in graphs, the concept of complete graphs is produced: every node is adjacent to every other node. Each production rule generates both a data table and a definition for the new concept, based on the data tables and definitions of the old concepts. HR uses seven production rules, each given in table 1 with a brief description of the types of concepts they produce.




Figure 1: re-invention of prime numbers

Figure 1 shows how the production rules can generate new concepts: prime numbers are constructed using just two production rules. Each production rule step is accompanied by a parameterisation detailing exactly how the step is to occur. For example, parameters $\langle 2 \rangle$ for the size production rule stipulate that for every object of interest, the number of sub-objects in column 2 is to be counted. The parameters $\langle 2 = 2 \rangle$ for the split production rule tell it to keep rows where column 2 contains value 2. Each production rule generates the set of all possible parameters for a given concept.

HR forms a theory using a heuristic search which builds new concepts from the best old ones. To decide which concepts are most worthy of development, HR has eight measures which calculate values based on the data table and construction history of the concepts. A weighted sum of the measures for each concept is used to order the concepts. One measure is the **complexity** of the concept, which determines how many concepts in total have been built upon to produce the new concept. Concepts with high complexity are often difficult to understand, so low complexity is usually favoured in the weighted sum. Often we impose a complexity limit of around 10 to improve the comprehensibility of the theory. The **novelty** measure of a concept calculates how many times the categorisation produced by the concept has been seen. For example, square numbers categorise integers into two sets: $[1, 4, 9, \dots]$ and $[2, 3, 5, \dots]$. If this categorisation had been seen often, square numbers would score poorly for novelty, and vice-versa.

To complete a cycle of mathematical activity, while searching for concepts, HR also checks for empirically true conjectures. For example, if a newly formed concept had exactly the same data table as a previous one, HR would make the conjecture that the definition for the old and new concepts were equivalent. Furthermore, when working in finite algebras, HR employs the Otter theorem prover (McCune 1990) to attempt to prove the conjectures it makes. HR then uses information from these proof attempts to re-assess the concepts in its theory - if a concept is involved in many interesting theorems, the interestingness of the concept is increased. For a more detailed discussion of the HR program, please see (Colton, Bundy, & Walsh 1999).

Extra Functionality for Integer Sequences

In number theory, we often start HR with only 3 concepts: integers - the objects of interest, divisors of integers - the sub-objects, and multiplication - relating two divisors if they multiply to give the integer. We can also supply other fundamental concepts, such as digits of integers and addition.

Name	Type of Concept Produced
Compose	identifying objects or sub-objects with properties from 2 concepts
Exists	identifying objects which have at least one sub-object with a particular property
forall	identifying objects for which all sub-objects have a particular property
Match	identifying objects with equal sub-objects satisfying a relation.
Negate	identifying objects or sub-objects which do not have a particular property.
Size	counts the number of sub-objects which have a particular property.
Split	identifying objects with a given number of sub-objects with a particular property.

Table 1: the seven production rules used by HR

There is also a choice of which integers to supply as the objects of interest - using too many will slow down the theory formation, but if we choose too few, many sequences which are different will appear the same. For instance, if we use the integers 1 to 5, the prime numbers, 2, 3, 5 will appear the same as the non-squares, which are also 2, 3 and 5. In practice, we use the numbers between 1 and around 15, depending on the initial concepts chosen.

We have enabled HR to present concepts in number theory as integer sequences. If the concept is a type of number, for example square numbers, HR just outputs the integers of that type in order. We call such sequences number type sequences. For concepts produced by the size production rule, eg. the τ -function, which counts the number of divisors of an integer, the sequence is formed from the values calculated by the function for the integers 1, 2, 3, etc. in order. We call these coefficient sequences. Other concepts may be special types of sub-objects, for instance prime divisors. To present these as sequences, HR takes each integer in turn and writes down the sub-objects in numerical order. We call such sequences sub-object sequences.

We have also added 3 production rules which take advantage of the natural ordering of integers, but which can work in any domain with a well defined ordering, eg. polynomials. These were inspired by the transformations used to identify sequences in the Encyclopedia (Bernstein & Sloane 1995) and we hope to implement more such rules. The first new production rule takes a coefficient sequence and finds those integers setting the 'record' for it. For example, highly composite numbers (sequence A002182) are those which set the record for the τ -function - they have more divisors than any smaller integer. The second production rule starts with a number type sequence and takes the difference between successive terms to produce a new sequence, a common way of producing new sequences. The third production rule specialises the 'extreme' production rule introduced in (Steel 1999). It looks at the sub-objects for each integer and finds either the largest or the smallest. This produces concepts such as the greatest prime divisor.

When using HR to invent sequences, we turn off the em-

pirical conjecture making and proving abilities, because, as discussed later, HR uses other methods to find conjectures. As long as the user supplies correct definitions for the initial concepts, HR will generate definitions consistent with the data for each concept. This therefore satisfies the criteria that sequences submitted to the Encyclopedia be well defined. HR can produce thousands of sequences, so there is the possibility that some simple sequences missing from the Encyclopedia can be identified and investigated.

Investigation of Sequences

When using HR to invent integer sequences, we set some parameters for the search and construct a fixed number of sequences, say 500. HR uses a local copy of the Encyclopedia to identify which of its inventions are missing, and orders them in terms of our complexity measure, so we can investigate the least complex ones first. We then examine the definitions of the sequences to choose a candidate for investigation and begin by using HR to calculate the terms up to 1000. Then HR compares the sequence to those already in the Encyclopedia and identifies any sequences which are related (as described below) with our new sequence. If HR finds such a relationship and we can prove it, we feel that the criteria for interestingness has been satisfied and we will probably submit the sequence to the Encyclopedia.

Unfortunately, the relationships identified can often be based on only little empirical evidence and may not be true in the general case. Also in many cases, there will be far too many sequences to which our sequence is related. For these reasons, we employ methods to prune the output, in an attempt to increase the yield of correct, interesting results. We first discuss the relationships that can be discovered, and then the pruning methods employed. As an example throughout, we use sequence A036436 which HR invented: integers with a square number of divisors. There are 62 terms for this stored in the Encyclopedia thus:

[A036436] 1, 6, 8, 10, 14, 15, 21, 22, 26, 27, . . . , 183

We first introduce some definitions. Note that these and all further definitions refer not to the idealised sequence (with an infinite number of terms), but rather to the sequences as they appear in the Encyclopedia, with a finite number of terms.

- The n th term of sequence S is written S_n , and we say that $a \in S$ if a is a term of S . The **number of terms** of S is written $|S|$.
- The **range** of a sequence, written $range(S)$, is the set of integers between the smallest term in the sequence and the largest term, inclusive.

For example, if S is sequence A036436 above, then $|S| = 62$ and $range(S) = \{1, 2, 3, 4, \dots, 183\}$.

Relationships Between Sequences

The following are three ways in which two sequences, S and T , can be related:

- S and T are **disjoint** if no term of S is a term of T .

- S is a **subsequence** of T if all the terms of S which are in $range(T)$ are also terms of T . Similarly, S is a **supersequence** of T if T is a subsequence of S .
- Letting k be the smaller of $|S|$ and $|T|$, we say that S is **less than** T if, for $i = 1, \dots, k, S_i \leq T_i$. Similarly, S is **greater than** T if T is less than S .

Once a relationship has been noted, the user must interpret the result as a mathematical conjecture. This is rarely difficult to do because of the simple nature of the relationships found. For example, HR notes that sequence A006881, integers of the form pq for distinct primes p and q , is a subsequence of A036436 above. This is interpreted as the following easy to prove fact: Any integer of the form pq has a square number of divisors. Similarly, HR notes that prime numbers are disjoint with A036436, which is interpreted as: no prime has a square number of divisors, also easy to prove.

Pruning Methods

When investigating sequence A036436 above, HR returns 3605 sequences from the Encyclopedia which are subsequences using the above definition. HR has many ways to prune the output, which are either constraints on the output sequence or constraints on both the sequence of interest and the output sequence. Firstly, we would like to discard sequences such as 1, 1, 1, 1, . . . which appear in the Encyclopedia for completeness, but are not particularly interesting. To do this, we can measure the number of *distinct* terms and discard any sequence with less than, say, 5 different terms.

If two sequences exist on different parts of the number line, it is uncertain whether one is a subsequence of the other. Our definition for subsequences admits those for which the range is disjoint with the range of the sequence being investigated. For example, HR notes that sequence A030091 (prime numbers, p , for which p and p^2 have the same digits) is a subsequence of A036436. This incorrect result occurs because the range of A030091 is $\{94583, \dots, 1029647\}$ which is disjoint with the range of A036436. We may want to discard sequences like this, which share no terms with our sequence. In general, we may wish to discard sequences which share only 1, 2, etc. terms with the new sequence. To do this, HR measures the number of shared terms:

- The **number of shared terms** of S and T is calculated as: $|\{a : a \in S \ \& \ a \in T\}|$.

As an example, sequence A036436 has terms 1, 36 and 100 stored in the Encyclopedia. These are the only square numbers in A036436, so the term overlap for this sequence with the sequence of square numbers is 3. This measure is very effective for pruning, eg. if we prune subsequences which share less than 3 terms with A036436, this reduces the number of results from 3605 to a more manageable 390. This measure is easier to use and as effective as a similar measure which determines the *proportion*, rather than the number of terms from one sequence that appear in another.

When looking for supersequences of a given sequence the, non-negative numbers: 0, 1, 2, 3, 4, . . . are always a supersequence. Therefore, we may want supersequences to be less dense on the number line than this, and HR uses the following measure for the density of a sequence:

- The **density** of S is calculated as:

$$\frac{|S|}{|\text{range}(S)|}$$

For example, there are 62 terms of A036436, distributed over the range of the first 183 integers. Therefore the density is $62/183 \approx 0.34$. Often we choose a limit which is only slightly larger than the density of the new sequence, as this can produce interesting results.

When looking for sequences which are disjoint with the sequence of interest, we want to avoid sequences where the ranges are disjoint, as the sequences are bound to be disjoint. Ideally the sequences should occupy roughly the same space on the number line (but without any overlapping terms). HR uses this measure for pruning trivially disjoint sequences:

- The **range overlap** of S and T is calculated as:

$$\frac{|\text{range}(S) \cap \text{range}(T)|}{|\text{range}(S) \cup \text{range}(T)|}$$

For example, in the Encyclopedia, the prime numbers have range $\{2, 3, 4, \dots, 271\}$ whereas sequence A036436 has range $\{1, 2, \dots, 183\}$. The range overlap of these sequences is therefore $\frac{|\{2,3,\dots,183\}|}{|\{2,3,\dots,271\}|} = \frac{182}{270} \approx 0.67$. If the minimum range overlap is set to close to 1, then the range of one sequence must be nearly contained in the range of the other. If this is true, yet the sequences are still disjoint, the result may be interesting.

When looking for a sequence which is less than our chosen sequence, it is desirable to look for sequences for which the terms are similar. HR calculates the average difference of the terms in the sequence thus:

- Letting k be the smaller of $|S|$ and $|T|$, the **difference** of S and T is calculated as:

$$\frac{1}{k} \sum_{i=1}^k |S_i - T_i|$$

If the maximum difference limit is set to 1, on average the n th terms of S and T will differ by only 1. Sequences so close to the new sequence may be interesting.

The final way to prune sequences is to use semantic information from the Encyclopedia. Firstly, HR can discard sequences with (or without) particular words in their definition. For example, when looking for subsequences of the prime numbers, it is desirable to discard any sequence with the word ‘prime’ in its definition, as these are usually obvious specialisations of primes, such as odd primes. Also, each sequence in the Encyclopedia has an associated set of keywords, such as ‘core’ - which are considered fundamental, and ‘nice’ - which have some appealing quality. HR can prune sequences if they have (or don’t have) particular keywords associated. For example, when we ask for subsequences of sequence A036436 which are described as ‘nice’ in the Encyclopedia, the output is reduced from 3605 sequences to just 34, one of which is A007422: integers n , where the product of the divisors of n is equal to n^2 . Like many of the results HR finds for A036436, a little investigation shows that this is true in the general case.

Results

Although we mainly use HR to invent new integer sequences, we would hope that it also re-invents many classically interesting sequences. We ran HR for 10 minutes, starting with the concepts of integers, divisors and multiplication. The search was depth limited to a complexity of 12, and we used the integers from 1 to 17. HR produced 233 sequences, of which 51 (22%) were sequences already in the Encyclopedia. This included 14 of the Encyclopedia’s 120 ‘core’ sequences with fundamental notions such as odd, even, square, cube and prime numbers being re-invented.

More complicated, non-core concepts were also re-invented, such as A000961: prime powers and A005117: square free numbers. Of the 51 re-invented sequences, 40 were found during the first minute, and so were less complex than the remaining 11 found in the final 9 minutes. This shows that less complex concepts are often more interesting, which is why we use a depth limited search. It also suggests more shallow searches using a variety of initial concepts rather than deep searches, where the complexity of the sequences makes them difficult to understand.

From all the number theory sessions so far, the total number of re-invented sequences is over 120. Certain sequences are re-invented with non-standard definitions, such as the sequence of powers of 2, which HR defines as: integers with only one odd divisor (eg. 1 is the only odd divisor of 1, 2, 4, 8, 16, etc. and all other integers have more odd divisors). A tactic which can increase the yield of classically interesting concepts is to identify the initial concepts and production rules required to re-invent a well known concept, and restrict the search to using only these. For example, the concept of the sum of divisors is reached using the concept of divisors and the “less than or equal to” concept. It is defined in this way, using only the size and compose production rules:

$$f(n) = |\{(d, a) : d|n \ \& \ a \leq d\}|.$$

Restricting theory formation to using only the concepts and production rules necessary for this construction, the first 14 sequences HR found were in the Encyclopedia. They included well known concepts such as the τ -function, square numbers, triangle numbers and the τ^2 -function.

HR re-invents many well known concepts because the production rules, while general in nature, were derived by studying the types of concepts found in mathematics. The compose rule is essential as it combines 2 concepts - without this, theory formation comes to a halt fairly quickly. Each re-invented concept required either the exists, split or size rule. These rules produce concepts identifying objects with certain sub-objects or a fixed number of sub-objects which are common constructions in number theory, eg. prime numbers, with 2 distinct divisors. The negate rule was also useful for constructing compliments of concepts, eg. odd numbers from even numbers. The number theory specific rules increased the yield of potentially interesting concepts with the ‘difference’ and ‘record’ rules effectively trebling the yield of sequences, as each new sequence was transformed by them into another one. The match and forall rules were less instrumental in this session, although both were required for at least one re-invented concept.

Of the 182 new sequences generated in the above 10 minute session, we assessed that 45 (25%) were unsuitable for investigation because either: (i) they were finite (eg. integers, n , where $\tau(n) = n$, the only examples being 1 and 2) or (ii) they were too specialised which made them dull (eg. the function $f(n) = |\{a : a|n \ \& \ n = 2 \times 2\}|$ produces a sequence of zeros except the fourth term which is 3). The remaining 137 sequences were worthy of investigation, and in fact 4 of them were sequences that HR previously invented which are now in the Encyclopedia (as discussed below). Judging from the 22% of sequences produced which were already in the Encyclopedia, HR is certainly generating the kinds of concepts which interest mathematicians, and that while some new sequences were not worthy of the Encyclopedia, 3/4 of those produced were candidates for investigation, which we feel is acceptable.

Illustrative Examples

To date, 17 sequences invented by HR have been added to the Encyclopedia. While HR invents hundreds of well defined sequences not present in the Encyclopedia, we have only submitted sequences for which HR has also found interesting conjectures. Every sequence we have submitted has been accepted, including:

[A036438] 1, 4, 6, 10, 12, 14, 22, 24, 26, 27, 32, 34, . . .

(integers expressible as $m \cdot \tau(m)$ for some m), and:

[A036433] 1, 2, 14, 23, 29, 34, 46, 63, 68, 74, 76, 78. . .

(integers for which the number of divisors is a digit).

The first of HR's sequences we submitted to the Encyclopedia was the refactorable numbers:

[A033950] 1, 2, 8, 9, 12, 18, 24, 36, 40, 56, 60, 72, . . .

which are those integers where the number of divisors is itself a divisor. We were informed later that these had been developed as recently as 1990 (Kennedy & Cooper 1990). On investigation, HR found 3 conjectures about refactorables which we have subsequently proved:

- Looking for disjoint sequences, and pruning with keywords, HR conjectured that perfect numbers [see (Sloane 2000)] are not refactorable.
- Looking for supersequences, HR conjectured that refactorables are only congruent to 0, 1, 2 or 4 mod 8.
- Looking for supersequences, HR conjectured that refactorables are of the form $lcm(a, \tau(a))$ for some a .

These, and more results found by us, were presented in a journal paper on refactorables (Colton 1999).

Investigation of sequences can be performed for any sequence, not just those invented by HR. For example, when looking for supersequences of fortunate numbers [see (Sloane 2000)], HR conjectured that they are all prime, a result known as Fortune's conjecture (Golomb 1981). We also used HR to look for supersequences of perfect numbers, and found that they are of the form $lcm(n, \sigma(n))$ for some n [where $\sigma(n)$ is the sum of the divisors of n]. This highlighted an appealing parallel between perfect numbers

and refactorables, which HR discovered were of the form $lcm(n, \tau(n))$, for some n .

When HR invented this sequence:

[A046952] 1, 4, 16, 36, 144, 576, 1296, 2304, 3600, . . .

which sets the record for this function:

$$f(n) = |\{(a, b) : a \times b = n \ \& \ a|b\}|, \quad (1)$$

it also made the conjecture that these are always square numbers. We went on to prove that the sequence was in fact the squares of the highly composite numbers.

Perhaps the most aesthetic result arose when HR invented the concept of integers where the number of divisors is prime, which is this sequence:

[A009087] 2, 3, 4, 5, 7, 9, 11, 13, 16, 17, 19, 23, 25, . . .

To investigate this we looked for subsequences described as 'nice' in the Encyclopedia. The first answer supplied was:

[A023194] 2, 4, 9, 16, 25, 64, 289, 729, 1681, 2401, . . .

which has the definition: integers for which the sum of divisors is prime. Therefore, HR made the conjecture that, given an integer, if the sum of divisors is prime, then the number of divisors will also be prime, which we subsequently proved. It is difficult to know whether this result is genuinely new, but it is certainly not well known, and is indicative of the kind of surprising and aesthetic conjectures it is possible to find using HR as an automated assistant.

Related Work

The aim of the SeekWhence program (Hofstadter 1995), (Meredith 1987) was not to invent sequences but to discover a definition for a given sequence. SeekWhence used heuristics to determine the nature of a sequence, such as taking the difference between two terms, or trying to extract and identify well known sub-sequences. For example, given the sequence 1, 1, 3, 4, 6, 9, SeekWhence would identify that square numbers: 1, 4, 9, and triangle numbers: 1, 3, 6, had been composed with repetition to form this sequence. Hofstadter aimed to model how humans search for reasonable definitions of sequences, rather than to provide a tool to identify sequences. The Guess program (Krattenthaler 1991) is such a tool which uses techniques from determinant calculus to produce a closed form definition for a given sequence. We have recently applied HR to the problem of identification and extrapolation of integer sequences (Colton, Bundy, & Walsh 2000), which requires more search control and is a distinct task to inventing sequences.

The AM program (Davis & Lenat 1982) worked in number theory, constructing a theory using a heuristic search to guide the invention of definitions. In contrast to HR, which starts with only a handful of concepts, 8 heuristic measures and 10 production rules (7, general, 3 specific to number theory), AM was supplied with 115 elementary concepts and used 242 heuristics to search for concepts. Some of these heuristics were very specific and often used only once during a session. AM re-invented well known sequences such as prime numbers and square numbers and the

τ -function and Lenat originally claimed that one of AM's sequences, integers with more divisors than any smaller integer was original. However, it later turned out that these had been defined as highly composite numbers and explored by Ramanujan (Hardy 1927). HR covers all the sequences re-invented by AM, and finds many outside of AM's range, such as powers of 2.

Conclusions and Further Work

The aim of the HR project is to provide a model for theory formation in pure mathematics and to investigate possible applications of this to mathematics and to areas of Artificial Intelligence such as machine learning. By implementing additional production rules and the ability to present concepts as sequences, we have applied HR to the invention of integer sequences. In a matter of minutes, it can re-invent more than 50 well known sequences, including 14 core sequences and can supply, in order of complexity, over 100 new sequences for investigation. We have linked HR to the Encyclopedia of Integer Sequences so the user can investigate a new sequence using HR to make conjectures about the sequence in relation to those in the Encyclopedia.

We have demonstrated that the theory formation techniques can scale up to produce interesting results and can be applied successfully in different domains. HR is the first program to both define new mathematical concepts and detail why they are of interest. This model for the invention of integer sequences has produced interesting novel results in number theory with the new sequences and conjectures generating genuine interest from mathematicians. For example, there are now over 30 sequences in the Encyclopedia related to refactorable numbers, submitted by various people.

The class of concepts which HR cannot invent is still large. In particular, it cannot invent concepts with recursive definitions, such as the factorial function. We are currently implementing a 'path' production rule which will output recursive definitions, thus increasing HR's coverage of these types of concepts. Also, many of the sequences in the Encyclopedia have aspects from more than one domain, eg. the first sequence in the Encyclopedia counts the number of groups with n elements. We are presently building on the work of (Steel 1999) to add cross domain functionality to HR. We have recently written the Integer Sequence Generator (ISG) as a Java implementation of the work presented here. We hope this will be an online tool which mathematicians use to invent and investigate new sequences. For instance, if the user was interested in refactorable numbers, the ISG would not only investigate refactorables, but would specialise the concept to, say, odd refactorable numbers. It would then conjecture, as we did in (Colton 1999), that all odd refactorables are squares.

Computer generated discoveries in mathematics such as the remarkable new formula for π detailed in (Bailey 1998) are still rare and often rely on specialised algorithms - in this case the PSLQ algorithm. Neil Sloane, the author of the Encyclopedia, states that coincidences arising from using the Encyclopedia have led to surprising conjectures (Sloane 1998). Whilst we cannot hope for such results as the famous moonshine conjectures (Conway & Norton 1979) which

arose coincidentally, we hope to further the possibility of surprising conjectures arising from the use of the Encyclopedia and theory formation programs such as HR.

Acknowledgements

This work is supported by EPSRC grant GR/M98012. We would like to thank the anonymous reviewers for their enthusiastic and helpful comments, and Neil Sloane for maintaining a truly remarkable encyclopedia.

References

- Bailey, D. 1998. Finding new mathematical identities via numerical computations. *ACM SIGNUM* 33(1):17 – 22.
- Bernstein, M., and Sloane, N. J. A. 1995. Some canonical sequences of integers. *Linear Algebra and its Applications* 226-228:57–72.
- Colton, S.; Bundy, A.; and Walsh, T. 1999. HR: Automatic concept formation in pure mathematics. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*.
- Colton, S.; Bundy, A.; and Walsh, T. 2000. Automatic identification of mathematical concepts. In *Proceedings of the 17th International Conference on Machine Learning*.
- Colton, S. 1999. Refactorable numbers - a machine invention. *Journal of Integer Sequences* 2.
- Conway, J., and Norton, S. 1979. Monstrous moonshine. *Bulletin of the London Mathematical Society* 11:308 – 339.
- Davis, R., and Lenat, D. 1982. *Knowledge-Based Systems in Artificial Intelligence*. McGraw-Hill Advanced Computer Science Series.
- Golomb, S. W. 1981. The evidence for Fortune's conjecture. *Mathematics Magazine* 54:209–210.
- Hardy, G. H., ed. 1927. *S. Ramanujan, Collected Papers*. Cambridge University Press.
- Hofstadter, D. 1995. *Fluid Concepts and Creative Analogies*. Basic Books.
- Kennedy, R., and Cooper, C. 1990. Tau numbers, natural density and Hardy and Wright's theorem 437. *International Journal of Mathematics and Mathematical Sciences* 13:2:383–386.
- Krattenthaler, C. 1991. Advanced determinant calculus. Technical report, Institute of Mathematics, University of Vienna.
- McCune, W. 1990. The OTTER user's guide. Technical Report ANL/90/9, Argonne National Laboratories.
- Meredith, M. J. E. 1987. *Seek-Whence: A Model of Pattern Perception*. Ph.D. Dissertation, Department of Computer Science, Indiana University.
- Sloane, N. J. A. 1998. My favorite integer sequences. In *Proceedings of the International Conference on Sequences and Applications*.
- Sloane, N. J. A. 2000. *Encyclopedia of Integer Sequences*. <http://www.research.att.com/~njas/sequences>.
- Steel, G. 1999. Cross domain theory formation. Master's thesis, Division of Informatics, University of Edinburgh.