Basic Concepts and Notation

Gabriel Robins

"When I use a word," Humpty Dumpty said, in a rather scornful tone, "it means just what I choose it to mean -- neither more nor less."

A **set** is formally an undefined term, but intuitively it is a (possibly empty) collection of arbitrary objects. A set is usually denoted by curly braces and some (optional) restrictions. Examples of sets are {1,2,3}, {hi, there}, and {k | k is a perfect square}. The symbol $\in$ denotes set **membership**, while the symbol $\notin$ denotes set **non-membership**; for example, $7 \in \{p \mid p \text{ prime}\}$ states that 7 is a prime number, while $q \notin \{0,2,4,6,...\}$ states that q is not an even number. Some **common sets** are denoted by special notation:

The **natural numbers**:	$\mathbb{N} = \{1,2,3,...\}$
The **integers**:	$\mathbb{Z} = \{...,-3,-2,-1,0,1,2,3,...\}$
The **rational numbers**:	$\mathbb{Q} = \{ \frac{a}{b} \mid a,b \in \mathbb{Z}, b \neq 0\}$
The **real numbers**:	$\mathbb{R} = \{x \mid x \text{ is a real number}\}$
The **empty set**:	$\emptyset = \{\}$

When only the **positive** elements of a numerical set are sought, a superscript "+" may be used to denote this. For example, $\mathbb{Z}^+ = \mathbb{N}$ denotes the positive integers (i.e., the natural numbers), $\mathbb{R}^+$ denotes all the positive reals, and more generally, $S^+ = \{s \in S \mid s>0\}$.

The logical symbol | (pronounced "such that", and sometimes also denoted as ∋) denotes a **conditional** (which usually follows this symbol). The logical symbol $\forall$ (pronounced "for all") denotes **universal quantification**. For example, the formula "$\forall x \in \mathbb{R}\ x \leq x^2+1$" reads "for all x a member of the real numbers, x is less than or equal to x-squared plus one" (i.e., no real number is greater than one more than its own square). The logical symbol $\exists$ (pronounced "there

exists") denotes **existential quantification**. For example, the formula "$\exists x \in \mathbb{Z} \mid x^2=5x$" states that there exists an integer whose square is equal to 5 times itself (i.e., x is either 5 or 0). These connectives may be composed in more complicated formulae, as in the following example: "$\forall x \in \mathbb{Z}\ \exists y \in \mathbb{Z} \mid y>x$" which states that there is no largest integer.

The logical connective $\wedge$ (pronounced "**and**") is a boolean-valued function that yields true if and only if both of its two arguments are true. The logical connective $\vee$ (pronounced "**or**") is a boolean-valued function that yields true if and only if one or more of its two arguments are true. The symbol $\Rightarrow$ (pronounced "**implies**") denotes logical implication; that is, $A \Rightarrow B$ means that B is true whenever A is true; for example, "$1 < x < y \Rightarrow x^3 < y^3$". The symbol $\Leftrightarrow$ (pronounced "if and only if", and sometimes written as "iff") denotes **logical equivalence**; that is, $A \Leftrightarrow B$ means that B is true if and only if A is true. More formally, $A \Leftrightarrow B$ means $A \Rightarrow B \wedge B \Rightarrow A$; an example is "$\min(x,y)=\max(x,y) \Leftrightarrow x=y$". It is easily shown that $A \Rightarrow B$ implies $\neg B \Rightarrow \neg A$, where $\neg$ denotes **logical negation**.

A set S is a **subset** of a set T (denoted $S \subseteq T$) if every element that is a member of S is also a member of T. More formally, $S \subseteq T \Leftrightarrow (x \in S \Rightarrow x \in T)$. A set S is a **proper** subset of a set T (denoted $S \subset T$) if S is a subset of T, but S and T are not equal. More formally, $S \subset T \Leftrightarrow (S \subseteq T \wedge S \neq T)$. Clearly every set has the empty set and the set itself as subsets (i.e., $\forall S\ \varnothing \subseteq S \wedge S \subseteq S$). Two sets are **equal** if and only if each is a subset of the other, i.e., $S=T \Leftrightarrow (T \subseteq S \wedge S \subseteq T)$.

The **union** of two sets S and T (denoted $S \cup T$) is the (duplicate-free) "merger" of the two sets. More formally, $S \cup T = \{x \mid x \in S \vee x \in T\}$. The **intersection** of two sets S and T (denoted $S \cap T$) is their greatest common subset. More formally, $S \cap T = \{x \mid x \in S \wedge x \in T\}$. Two sets are said to be disjoint if their intersection is empty (i.e., S and T are disjoint $\Leftrightarrow S \cap T = \varnothing$).

The union and intersection operators are **commutative** ($S\cup T=T\cup S$, and $S\cap T=T\cap S$), **associative** $S\cup(T\cup V) = (S\cup T)\cup V$, and $S\cap(T\cap V) = (S\cap T)\cap V$, and **distribute** over each other $S\cup(T\cap V)=(S\cup T)\cap(S\cup V)$, and $S\cap(T\cup V) = (S\cap T)\cup(S\cap V)$. **Absorption** occurs as follows: $S\cup(S\cap T)=S$, and $S\cap(S\cup T)=S$. The **complement** of a set S (with respect to some universe set) is the collection of all elements (in the universe set) that are not in S, and is denoted S' (or by S with a horizontal bar over it). More formally, $S' = \{x \mid x\notin S\}$.

A set is said to be **closed** under a given operation if the operation preserves membership in the set. Formally, S is said to be closed under an operation $\Diamond$ iff $x\Diamond y\in S\ \forall\ x,y\in S$. For example, the set of integers $\mathbb{Z}$ is closed under addition (+), since the sum of any two integers is also an integer; on the other hand, $\mathbb{Z}$ is not closed under division.

A **relation** over a domain D is a set of ordered pairs, or more generally, a set of ordered k-tuples. For example, the relation $\heartsuit$ defined as {(a,1), (b,2), (b,3)} means that "a" is related to 1, and "b" is related to both 2 and 3; this may also be written as $a\heartsuit 1$, $b\heartsuit 2$, and $b\heartsuit 3$. A more familiar relation (over $\mathbb{Z}$) is the "less than" relation, often denoted as <, which actually consists of an infinite set of ordered pairs such that the first element is less than the second; that is, the < relation is formally defined to be the set $\{(x,y) \mid x,y\in\mathbb{Z}, y>x\}$.

A relation is said to be **reflexive** if every element in the relation domain is also related to itself; i.e., $\heartsuit$ is reflexive iff $x\heartsuit x\ \forall x\in D$. A relation is said to be **symmetric** if it commutes; i.e., $\heartsuit$ is symmetric iff $x\heartsuit y \Rightarrow y\heartsuit x$. A relation is **transitive** if $x\heartsuit y \wedge y\heartsuit z \Rightarrow x\heartsuit z$. For example, the subset operator is **reflexive** ($S\subseteq S$), and **transitive** ($S\subseteq T \wedge T\subseteq V \Rightarrow S\subseteq V$), but not symmetric. The **transitive closure** of a relation is the extension of that relation to all pairs that are related by transitivity; i.e., the transitive closure of $\heartsuit$ contains all pairs of $\heartsuit$, as well as all pairs (x,y) such that for some finite set of elements $d_1, d_2, d_3, \ldots, d_k$ in $\heartsuit$'s domain, all of $x\heartsuit d_1$, $d_1\heartsuit d_2$, $d_2\heartsuit d_3$, ..., $d_{k-1}\heartsuit d_k$, $d_k\heartsuit y$ hold. Put another way, the transitive closure $\spadesuit$ of a relation $\heartsuit$

is the smallest relation containing ♥ but which is still closed under transitivity (i.e., satisfying $x \spadesuit y \wedge y \spadesuit z \Rightarrow x \spadesuit z$). For example, the predecessor relation ‡ may be defined as $\{(x,x-1) \mid x \in \mathbb{Z}\}$, and the transitive closure of ‡ is the > relation. Similarly, the **symmetric closure** of a relation is the smallest containing relation that is closed under symmetry, etc.

"I don't understand you," said Alice. "Its dreadfully confusing!"

A relation that is reflexive, symmetric, and transitive is called an **equivalence relation**; an example of this is the familiar equality relation =. It is easy to show that an equivalence relation partitions its domain into mutually disjoint subsets, called **equivalence classes**. A special kind of relation is called a **graph**, where the domain elements are called **nodes** and the relation pairs are referred to as **edges**. For example, one simple graph may be {(a,b), (a,c), (b,d)}. Graphs are often drawn using ovals to represent the nodes and arcs to represent the edges. A graph is said to be **undirected** when the relation that it represents is symmetric, and **directed** otherwise. The transitive closure of an undirected graph is an equivalence relation where the equivalence classes correspond to the connected components of the graph.

"You'll get used to it in time," said the Caterpillar;

An important property of set operations is the analogue of **DeMorgan's Law**: $(S \cup T)' = S' \cap T'$ and $(S \cap T)' = S' \cup T'$. These equalities follow from DeMorgan's law for classical logic: if X and Y are boolean variables, then $(X \wedge Y)' = X' \vee Y'$ and $(X \vee Y)' = X' \wedge Y'$ always hold. This is an artifact of the elegant duality between the operators $\wedge$ and $\vee$ in the prepositional calculus: if one starts with a true theorem (logical preposition) and simultaneously replaces all the $\wedge$'s with $\vee$'s, and all the $\vee$'s with $\wedge$'s, the result is also a true theorem.

The **difference** between two sets S and T is the set containing all elements that are in S but not in T. More formally, $S-T = \{s \mid s \in S \wedge s \notin T\} = S \cap T'$. The **symmetric difference** between two sets S and T is defined as $S \cup T - S \cap T$. The **cross-product** of two sets S and T, denoted by $S \times T$, is the set of all ordered pairs whose first element comes from S and whose second element

comes from T. More formally, $S \times T = \{(s,t) \mid s \in S, t \in T\}$. For example, $\{1, 2, 3\} \times \{a,b\} = \{(1,a),(1,b),(2,a),(2,b),(3,a),(3,b)\}$. A set may be crossed with itself a number of times: $S^i = S \times S^{i-1}$ where $S^1 = S$.

The **cardinality** (or size) of a finite set is defined to be the number of elements in it, and is denoted by vertical bars placed around the set. For example, $|\{a,b,c\}|=3$, $|\{p \mid p \text{ a prime less than } 20\}|=8$, and $|\emptyset|=0$. The **powerset** of a set S (denoted 2^S) is the collection of all subsets of S; more formally, $2^S=\{T \mid T \subseteq S\}$. If S is finite, the cardinality of its powerset is precisely 2 raised to the cardinality of S (i.e., $|2^S|=2^{|S|}$); this is true because each subset of S can be represented by a unique sequence of |S| binary digits (where 1 represents membership of the corresponding element in the subset, and 0 represents non-membership). Since there are $2^{|S|}$ such sequences, and each corresponds to a unique subset of S, there must be $2^{|S|}$ subsets of S.

A function f which maps a set S to a set T (denoted $f:S \rightarrow T$) is said to be **one-to-one** (or **injective**) if any two distinct elements in S are mapped to distinct elements in T. More formally, f is injective iff $a,b \in S \therefore a \neq b \Rightarrow f(a) \neq f(b)$. In this context S is said to be the **domain** of f, while T is said to be the **range** of f. Intuitively, a function is one-to-one if no two distinct elements in its domain are mapped to the same element in its range. For example, $f:\mathbb{Z} \rightarrow \mathbb{Z}$ defined as $f(x)=2x$ is one-to-one, while $g(x)=x^2$ is not, since g maps both -2 and 2 to 4.

The **rate of growth** of numerical functions is often described asymptotically. This is especially useful when discussing the time or space complexities of algorithms, since it enables implementation- / hardware-independent comparisons of the relative merits of specific algorithms. A function $f(x)$ is said to be **O(g(x))** (pronounced "**big-oh** of g(x)") if for some positive constant c we have $c \cdot f(x) < g(x)$ for all but a finite number of values of x. If other words, g(x) is an **upper bound** for $f(x)$ in the limit, modulo a multiplicative constant. More formally, this may be expressed as $\boldsymbol{f(x) = O(g(x))} \Leftrightarrow \exists\, c \in \mathbb{R}^+ \exists\, x' \in \mathbb{R}^+ \ni f(x) < c \cdot g(x) \ \forall\, x > x'$. Similarly, a

function $f(x)$ is said to be $\Omega(g(x))$ (pronounced "**omega** of g(x)") if for some positive constant c we have $c \cdot f(x) > g(x)$ for all but a finite number of values of x. If other words, g(x) is a **lower bound** for $f(x)$ in the limit, modulo a multiplicative constant. More formally, this may be expressed as $f(x) = \Omega(g(x)) \Leftrightarrow \exists\, c \in \mathbb{R}^+ \;\exists\, x' \in \mathbb{R}^+ \ni f(x) > c \cdot g(x)\ \forall\, x > x'$.

Finally, a function $f(x)$ is said to be $\Theta(g(x))$ (pronounced "**theta** of g(x)") if both the relations $f(x) = \Omega(g(x))$ and $f(x) = O(g(x))$ hold; in other words, $f(x)$ and g(x) have the same asymptotic growth rate, modulo a multiplicative constant, so that each of $f(x)$ and g(x) gives a **tight bound** (or exact bound) for the other. For example, $f(n) = n$ is both O(n) and also $O(n^3)$. Similarly, $g(n) = 8 \cdot n \log n$ is $\Omega(n)$ and $O(n^{1.5})$, but not $\Omega(n^2)$. Finally, the constant function $h(n) = 100^{100}$ is O(1), as is any constant, no matter how large. Note that care must be taken when considering asymptotic notation; for example, h(x) = O(1) does not imply that h is a constant function, since non-constant yet bounded functions such as h(x) = sin(x) are also O(1) by the above definitions. Both O and Ω are reflexive and transitive relations, but are not commutative. On the other hand, is Θ is commutative as well.

$f: S \rightarrow T$ is said to be **onto** (or **surjective**) if for any element t in T, there exists an element s in S such that $f(s)=t$. More formally, f is onto iff $\forall\, t \in T\ \exists\, s \in S \ni f(s)=t$. Intuitively, a function is onto if its entire range is "covered" by its domain. For example, $f: \mathbb{Z} \rightarrow \mathbb{Z}$ defined as f(x)=13-x is onto (and coincidentally one-to-one as well), while $g(x)=x^2$ is not, since some elements of g's range do not have a corresponding element x in g's domain (i.e., there is no integer k such that g(k)=3).

A function that is both injective and surjective is called **bijective**, and is said to be (or to constitute) a **one-to-one-correspondence** between its domain and range. Intuitively, a bijection (denoted $\leftrightarrow$) is a perfect pairwise matching between two sets, with each element in each set participating in *exactly* one match with an element of the other set. For example, the identity

function on an arbitrary domain D is always a bijection (i.e. f:D $\leftrightarrow$D $\ni$ f(x)=x). Another example of a bijection is h:$\mathbb{N}\rightarrow\mathbb{Z}$ defined as h(x)= $\frac{x-1}{2}$ if x is odd, $\frac{-x}{2}$ if x is even. The last example illustrates the fact that an infinite set can be put into one-to-one correspondence with a *proper* subset of itself! (which is of course *never* possible for a finite set).

The cardinality of a set S is said to be **at least as large** as the cardinality of a set T, if there exists an onto function from S to T. More formally, $|S|\geq|T| \Leftrightarrow \exists\ f:S\rightarrow T$, f is onto. Note how this definition generalizes the notion of cardinality comparisons to infinite sets. For example, the onto function r:$\mathbb{R}\rightarrow\mathbb{Z}$ defined as "r(x)=integer closest to x" is witness to the fact that the reals have a cardinality at least as large as the integers.

If $|S|\geq|T|$ and a bijection between S and T exists, the cardinality of S is said to be **the same** as the cardinality of T. If $|S|\geq|T|$ but no bijection between S and T exists, the cardinality of S is said to be **strictly larger** than the cardinality of T, denoted $|S|>|T|$. The bijection h defined earlier proves that the natural numbers have the same cardinality as do the integers, even though the former is a proper subset of the latter!

It turns out that the cardinality of the reals is strictly larger than the cardinality of the natural numbers (formally $|\mathbb{R}|>|\mathbb{N}|$). This can be proved using a **diagonalization** argument: we already know that $|\mathbb{R}|\geq|\mathbb{N}|$, since y:$\mathbb{R}\rightarrow\mathbb{N}$ defined as "y(x)=abs(truncate(x))" is onto. Now assume that there exists an arbitrary bijection f:$\mathbb{N}\leftrightarrow\mathbb{R}$. Now consider the real number Ω defined so that Ω's $k^{\underline{th}}$ digit (to the right of the decimal point) is equal to [f(k)'s $k^{\underline{th}}$ digit] + 1 (modulo 10), for k=1,2,3,... Clearly Ω is a well-defined real number, but is *not* in the range of f by construction. It follows that f therefore cannot be a bijection as claimed, and since f was arbitrary, no bijection between $\mathbb{R}$ and $\mathbb{N}$ can possibly exist. Diagonalization is a powerful proof method which is often employed to establish non-existence results.

Bijections may be composed to form new bijections, so that if we have two bijections $a:S\rightarrow T$ and $b:T\rightarrow V$, then we can form a new bijection $c:S\rightarrow V$, defined as $c(x)=b(a(x))$. As an example of an application of this **composition principle**, we note that no bijection between $\mathbb{R}$ and $\mathbb{Z}$ can possibly exist: h (as defined earlier) is a bijection between $\mathbb{N}$ and $\mathbb{Z}$, and we already know that no bijection between $\mathbb{R}$ and $\mathbb{N}$ can possibly exist (by our earlier diagonalization proof). Therefore a bijection between $\mathbb{R}$ and $\mathbb{Z}$ would automatically yield (using our composition principle) a bijection between $\mathbb{R}$ and $\mathbb{N}$, a contradiction.

"Oh!" said Alice. She was too much puzzled to make any other remark.

An **infinite set** is a set than can be put into one-to-one correspondence with a *proper* subset of itself (or intuitively, a set with a cardinality greater than any integer $k\in\mathbb{Z}$). Any set that is finite, or else that can be put into one-to-one correspondence with the integers is said to be **countable** (or countably infinite). Any infinite set that *can not* be put into one-to-one correspondence with the integers is said to be **uncountable** (or uncountably infinite). For example, $\mathbb{N}$, $\mathbb{Z}\times\mathbb{Z}$, $\mathbb{Q}$, and $\{p \mid p \text{ prime}\}$ are all countable sets, while $\mathbb{R}$, $\{x \mid x\in\mathbb{R}, 0\leq x\leq 1\}$, and $2^{\mathbb{N}}$ are all uncountable sets.

An **alphabet** is a finite set of symbols (e.g., $\Sigma = \{a,b,c\}$). A **string** is a finite sequence of symbols chosen from a particular alphabet (e.g., w = abcaabbcc). The **length** of a string is the numbers of symbols it is composed of (e.g., $|bca| = 3$). A **language** is a set of strings over some alphabet. For example, for the alphabet $\Sigma=\{a,b\}$, aaabbbabab is a string of length 10 over Σ, and $\{a^nb^n \mid n>0\}$ is an infinite language over Σ. The unique string of length 0 is the **empty string**, and is denoted by ε or ^. The **concatenation** of two strings x and y (denoted xy) is obtained by following the symbols of x with the symbols of y, in order. More formally, if $x=x_1x_2x_3\ldots x_n$ and $y=y_1y_2y_3\ldots y_m$, where $x_i\in\Sigma$ for $1\leq i\leq n$ and $y_j\in\Sigma$ for $1\leq j\leq m$, then $xy=x_1x_2x_3\ldots x_ny_1y_2y_3\ldots y_m$. It follows that for all strings w over some alphabet, $w\varepsilon=\varepsilon w=w$. For example, the string "hi" concatenated to the string "there" yields the string "hithere".

The concatenation operator may be extended to languages L_1 and L_2 as follows: $L_1L_2 = \{xy \mid x \in L_1 \text{ and } y \in L_2\}$. LL may be denoted by L^2; more generally, $L^k = LL^{k-1}$, where $L^0 = \{\varepsilon\}$. The **Kleene closure** of a language L (denoted by L^*) is defined as the infinite union $L^0 \cup L^1 \cup L^2 \cup L^3 \cup \ldots$, while L^+ is defined as the infinite union $L^1 \cup L^2 \cup L^3 \cup \ldots$ It follows that $L^+ = LL^*$ (note that this "superscript plus" notation is distinguished from the "superscript plus" used earlier to denote the positive elements of a numerical set, e.g., $\mathbb{Z}^+ = \mathbb{N}$, and usually the context may be consulted to avoid confusion).

For example, the language {a, b} concatenated to the language {1, 2, 3} yields the language {a1, a2, a3, b1, b2, b3}, while $\{a,b\}^*$ denotes the set of all finite strings over the two symbols a and b; more generally, Σ^* denotes the set of all finite strings over the alphabet Σ. It turns out that $(L^*)^* = L^*$, and that unless L is the **trivial language** (i.e., $\{\varepsilon\}$) or the **empty language** (i.e., $\emptyset$) then L^* is countably infinite. Note that the trivial language $\{\varepsilon\}$ is not the same as the empty language $\emptyset$: the former contains one exactly string (i.e., the empty string) but the latter contains none.

Any language L over a finite alphabet Σ is composed of some collection of finite strings. More formally, $L \subseteq \Sigma^*$. Clearly Σ^* is countable (simply arrange the finite strings in Σ^* by increasing length, and within length by lexicographic dictionary order). Similarly, the set of all **finite descriptions** is countable (simply arrange the description by increasing lengths and lexicographically within the same length). On the other hand, the set of all languages 2^{Σ^*} is uncountable. This immediately implies that some languages are not finitely describable! Put differently, the set of all possible finite algorithms (or descriptions) is countable (sort the finite computer programs by size and lexicographically), while the set of problems (languages) is uncountable; this means that any way of matching solutions to problems must leave out some problems unmatched (i.e., unsolved), and therefore some problems have absolutely no solutions, even in theory! Exhibiting an actual "finitely undescribable" set requires a little more work, but is not altogether difficult; this is what Alan Turing did in his 1936 dissertation.

"Curiouser and curiouser!" cried Alice.

The infinity corresponding to the **cardinality of the integers** is denoted by $\aleph_0$ (pronounced "aleph null"). The infinity corresponding to the **cardinality of the reals** is denoted as $\aleph_1$. Our previous discussion established that $\aleph_0 < \aleph_1$, and formally we have $\aleph_1 = 2^{\aleph_0}$. For many years mathematicians have tried to find some infinity Ω such that $\aleph_0 < \Omega < \aleph_1$, or prove that none exists. This question of whether there exists some infinity strictly larger than that of the integers, yet strictly smaller than that of the reals, came to be known as the "**continuum hypothesis**," and was finally settled by Cohen in 1966, who showed that to be **independent of the axioms** of set theory. That is, the consistency of set theory would not be changed if one chooses to assume as an axiom either this hypothesis, *or* its negation! Several other well-known mathematical statements enjoy this unique status of being independent of the axioms, and these include the **parallel postulate**, as well as the **axiom of choice** (shown to be independent of the other axioms by Godel in 1938).

More generally, we can obtain a whole **hierarchy of infinities**, each one strictly greater than its predecessor; in particular, we have $\aleph_{i+1} = 2^{\aleph_i}$, where $\aleph_i < \aleph_{i+1}$. But when the indexes of the alephs keep growing, nothing prevents them from soon becoming alephs themselves! In other words, our "number-line" now looks like:

$$0,\ 1,\ 2,...,\ k,\ k+1,...,\ \aleph_0\ ,\ \aleph_1,\ \aleph_2\ ,...,\ \aleph_k,\ \aleph_{k+1}\ ,...,\ \aleph_{\aleph_0},\ \aleph_{\aleph_1},...,\ \aleph_{\aleph_k},\ \aleph_{\aleph_{k+1}},\ldots$$

where the subscripts soon acquire subscripts which are themselves alephs, giving rise to an infinite hierarchy of infinities! Does there exist any infinity "bigger" then any of these unimaginably large cardinalities? It turns out that there is! The next "jump" in this sequence is denoted by ω (pronounced "**omega**") and is bigger than any of the alephs "below" it. It is sometimes referred to as the "**first inaccessible infinity**" because there is no way to "reach" it via any composition,

exponentiation, or subscript-nesting of alephs, etc., very much like there is no way to reach the first aleph via any finite sequence of arithmetic operations on the ordinary integers.

> *The Red Queen shook her head. "you may call it 'nonsense' if you like," she said, "but I've heard nonsense, compared with which that would be as sensible as a dictionary!"*

Interestingly, this fascinating progression of ever-increasing infinities does not stop; using certain logical constructs is it possible to exhibit a vast hierarchy of inaccessible infinities past ω! Logicians have even "found" infinities "larger" than any of the inaccessible ones, by stretching the power of their axiomatic proof systems to the limit. Note that finding a new families of infinities requires new and novel proof techniques, since the "jump" from one "level" of infinities to the next "level" is as fundamental and conceptually difficult as the initial jump from the integers to the first level at $\aleph_0$, or the jump from the alephs to ω! Currently only about six more fundamental "jumps" in conceptualization are known to logicians, enjoying names such as the hyper-Mahlo cardinals, the weakly compact cardinals, and the ineffable cardinals. It is not clear (even in theory) what exotic mathematical constructs, if any, lay beyond that.

References

Behnke, H., Bachmann, Fladt, K., and Suss, W., (eds.), Fundamentals of Mathematics, Volume I, MIT Press, Cambridge, Mass, 1973, pp. 50-61.

Davis, P., Hersh, R., The Mathematical Experience, Birkhauser, Boston, 1980, pp.152-157, 223-236.

Harel, D., Algorithmics: the Spirit of Computing, 2nd Ed., Addison-Wesley, 1992.

Hilbert, D., On the Infinite, in Benacerrat, P., and Putnam, H., (eds.), Philosophy of Mathematics: Selected Readings, Englewood Cliffs: Prentice-Hall, 1964, pp. 134-151.

Hopcroft, J., and Ullman, J., Introduction to Automata Theory, Languages, and Computation, Addison-Wesley, Reading, Massachusetts, 1979.

Rucker, R., Infinity and the Mind: the Science and Philosophy of the Infinite, Harvester Press, 1982.

Zippin, L., Uses of Infinity, Mathematical Association of America, Washington, D.C., 1962.