

udf Induction

ind.1 Introduction

mth:ind:int:
sec Induction is an important proof technique which is used, in different forms, in almost all areas of logic, theoretical computer science, and mathematics. It is needed to prove many of the results in logic.

Induction is often contrasted with deduction, and characterized as the inference from the particular to the general. For instance, if we observe many green emeralds, and nothing that we would call an emerald that's not green, we might conclude that all emeralds are green. This is an inductive inference, in that it proceeds from many particular cases (this emerald is green, that emerald is green, etc.) to a general claim (all emeralds are green). *Mathematical induction* is also an inference that concludes a general claim, but it is of a very different kind than this “simple induction.”

Very roughly, an inductive proof in mathematics concludes that all mathematical objects of a certain sort have a certain property. In the simplest case, the mathematical objects an inductive proof is concerned with are natural numbers. In that case an inductive proof is used to establish that all natural numbers have some property, and it does this by showing that

1. 0 has the property, and
2. whenever a number k has the property, so does $k + 1$.

Induction on natural numbers can then also often be used to prove general claims about mathematical objects that can be assigned numbers. For instance, finite sets each have a finite number n of elements, and if we can use induction to show that every number n has the property “all finite sets of size n are . . .” then we will have shown something about all finite sets.

Induction can also be generalized to mathematical objects that are *inductively defined*. For instance, expressions of a formal language such as those of first-order logic are defined inductively. *Structural induction* is a way to prove results about all such expressions. Structural induction, in particular, is very useful—and widely used—in logic.

ind.2 Induction on \mathbb{N}

mth:ind:inN:
sec In its simplest form, induction is a technique used to prove results for all natural numbers. It uses the fact that by starting from 0 and repeatedly adding 1 we eventually reach every natural number. So to prove that something is true for every number, we can (1) establish that it is true for 0 and (2) show that whenever it is true for a number n , it is also true for the next number $n + 1$. If we abbreviate “number n has property P ” by $P(n)$ (and “number k has property P ” by $P(k)$, etc.), then a proof by induction that $P(n)$ for all $n \in \mathbb{N}$ consists of:

1. a proof of $P(0)$, and
2. a proof that, for any k , if $P(k)$ then $P(k + 1)$.

To make this crystal clear, suppose we have both (1) and (2). Then (1) tells us that $P(0)$ is true. If we also have (2), we know in particular that if $P(0)$ then $P(0 + 1)$, i.e., $P(1)$. This follows from the general statement “for any k , if $P(k)$ then $P(k + 1)$ ” by putting 0 for k . So by modus ponens, we have that $P(1)$. From (2) again, now taking 1 for n , we have: if $P(1)$ then $P(2)$. Since we’ve just established $P(1)$, by modus ponens, we have $P(2)$. And so on. For any number n , after doing this n times, we eventually arrive at $P(n)$. So (1) and (2) together establish $P(n)$ for any $n \in \mathbb{N}$.

Let’s look at an example. Suppose we want to find out how many different sums we can throw with n dice. Although it might seem silly, let’s start with 0 dice. If you have no dice there’s only one possible sum you can “throw”: no dots at all, which sums to 0. So the number of different possible throws is 1. If you have only one die, i.e., $n = 1$, there are six possible values, 1 through 6. With two dice, we can throw any sum from 2 through 12, that’s 11 possibilities. With three dice, we can throw any number from 3 to 18, i.e., 16 different possibilities. 1, 6, 11, 16: looks like a pattern: maybe the answer is $5n + 1$? Of course, $5n + 1$ is the maximum possible, because there are only $5n + 1$ numbers between n , the lowest value you can throw with n dice (all 1’s) and $6n$, the highest you can throw (all 6’s).

Theorem ind.1. *With n dice one can throw all $5n + 1$ possible values between n and $6n$.*

Proof. Let $P(n)$ be the claim: “It is possible to throw any number between n and $6n$ using n dice.” To use induction, we prove:

1. The *induction basis* $P(1)$, i.e., with just one die, you can throw any number between 1 and 6.
2. The *induction step*, for all k , if $P(k)$ then $P(k + 1)$.

(1) Is proved by inspecting a 6-sided die. It has all 6 sides, and every number between 1 and 6 shows up one on of the sides. So it is possible to throw any number between 1 and 6 using a single die.

To prove (2), we assume the antecedent of the conditional, i.e., $P(k)$. This assumption is called the *inductive hypothesis*. We use it to prove $P(k + 1)$. The hard part is to find a way of thinking about the possible values of a throw of $k + 1$ dice in terms of the possible values of throws of k dice plus of throws of the extra $k + 1$ -st die—this is what we have to do, though, if we want to use the inductive hypothesis.

The inductive hypothesis says we can get any number between k and $6k$ using k dice. If we throw a 1 with our $(k + 1)$ -st die, this adds 1 to the total. So we can throw any value between $k + 1$ and $6k + 1$ by throwing k dice and then rolling a 1 with the $(k + 1)$ -st die. What’s left? The values $6k + 2$ through

$6k + 6$. We can get these by rolling k 6s and then a number between 2 and 6 with our $(k + 1)$ -st die. Together, this means that with $k + 1$ dice we can throw any of the numbers between $k + 1$ and $6(k + 1)$, i.e., we've proved $P(k + 1)$ using the assumption $P(k)$, the inductive hypothesis. \square

Very often we use induction when we want to prove something about a series of objects (numbers, sets, etc.) that is itself defined "inductively," i.e., by defining the $(n + 1)$ -st object in terms of the n -th. For instance, we can define the sum s_n of the natural numbers up to n by

$$\begin{aligned} s_0 &= 0 \\ s_{n+1} &= s_n + (n + 1) \end{aligned}$$

This definition gives:

$$\begin{aligned} s_0 &= 0, \\ s_1 &= s_0 + 1 &&= 1, \\ s_2 &= s_1 + 2 &&= 1 + 2 = 3 \\ s_3 &= s_2 + 3 &&= 1 + 2 + 3 = 6, \text{ etc.} \end{aligned}$$

Now we can prove, by induction, that $s_n = n(n + 1)/2$.

Proposition ind.2. $s_n = n(n + 1)/2$.

Proof. We have to prove (1) that $s_0 = 0 \cdot (0 + 1)/2$ and (2) if $s_k = k(k + 1)/2$ then $s_{k+1} = (k + 1)(k + 2)/2$. (1) is obvious. To prove (2), we assume the inductive hypothesis: $s_k = k(k + 1)/2$. Using it, we have to show that $s_{k+1} = (k + 1)(k + 2)/2$.

What is s_{k+1} ? By the definition, $s_{k+1} = s_k + (k + 1)$. By inductive hypothesis, $s_k = k(k + 1)/2$. We can substitute this into the previous equation, and then just need a bit of arithmetic of fractions:

$$\begin{aligned} s_{k+1} &= \frac{k(k + 1)}{2} + (k + 1) = \\ &= \frac{k(k + 1)}{2} + \frac{2(k + 1)}{2} = \\ &= \frac{k(k + 1) + 2(k + 1)}{2} = \\ &= \frac{(k + 2)(k + 1)}{2}. \end{aligned} \quad \square$$

The important lesson here is that if you're proving something about some inductively defined sequence a_n , induction is the obvious way to go. And even if it isn't (as in the case of the possibilities of dice throws), you can use induction if you can somehow relate the case for $k + 1$ to the case for k .

ind.3 Strong Induction

In the principle of induction discussed above, we prove $P(0)$ and also if $P(k)$, then $P(k+1)$. In the second part, we assume that $P(k)$ is true and use this assumption to prove $P(k+1)$. Equivalently, of course, we could assume $P(k-1)$ and use it to prove $P(k)$ —the important part is that we be able to carry out the inference from any number to its successor; that we can prove the claim in question for any number under the assumption it holds for its predecessor.

There is a variant of the principle of induction in which we don't just assume that the claim holds for the predecessor $k-1$ of k , but for all numbers smaller than k , and use this assumption to establish the claim for k . This also gives us the claim $P(n)$ for all $n \in \mathbb{N}$. For once we have established $P(0)$, we have thereby established that P holds for all numbers less than 1. And if we know that if $P(l)$ for all $l < k$, then $P(k)$, we know this in particular for $k=1$. So we can conclude $P(1)$. With this we have proved $P(0)$ and $P(1)$, i.e., $P(l)$ for all $l < 2$, and since we have also the conditional, if $P(l)$ for all $l < 2$, then $P(2)$, we can conclude $P(2)$, and so on.

In fact, if we can establish the general conditional “for all k , if $P(l)$ for all $l < k$, then $P(k)$,” we do not have to establish $P(0)$ anymore, since it follows from it. For remember that a general claim like “for all $l < k$, $P(l)$ ” is true if there are no $l < k$. This is a case of vacuous quantification: “all A s are B s” is true if there are no A s, $\forall x (\varphi(x) \rightarrow \psi(x))$ is true if no x satisfies $\varphi(x)$. In this case, the formalized version would be “ $\forall l (l < k \rightarrow P(l))$ ”—and that is true if there are no $l < k$. And if $k=0$ that's exactly the case: no $l < 0$, hence “for all $l < 0$, $P(0)$ ” is true, whatever P is. A proof of “if $P(l)$ for all $l < k$, then $P(k)$ ” thus automatically establishes $P(0)$.

This variant is useful if establishing the claim for k can't be made to just rely on the claim for $k-1$ but may require the assumption that it is true for one or more $l < k$.

ind.4 Inductive Definitions

In logic we very often define kinds of objects *inductively*, i.e., by specifying rules for what counts as an object of the kind to be defined which explain how to get new objects of that kind from old objects of that kind. For instance, we often define special kinds of sequences of symbols, such as the terms and formulas of a language, by induction. For a simple example, consider strings consisting of letters a, b, c, d, the symbol \circ , and brackets [and], such as “[c \circ d]”, “[a[] \circ]”, “a” or “[a \circ b] \circ d]”. You probably feel that there's something “wrong” with the first two strings: the brackets don't “balance” at all in the first, and you might feel that the “ \circ ” should “connect” expressions that themselves make sense. The third and fourth string look better: for every “[” there's a closing “]” (if there are any at all), and for any \circ we can find “nice” expressions on either side, surrounded by a pair of parentheses.

We would like to precisely specify what counts as a “nice term.” First of all,

every letter by itself is nice. Anything that's not just a letter by itself should be of the form " $[t \circ s]$ " where s and t are themselves nice. Conversely, if t and s are nice, then we can form a new nice term by putting a \circ between them and surround them by a pair of brackets. We might use these operations to *define* the set of nice terms. This is an *inductive definition*.

Definition ind.3 (Nice terms). The set of *nice terms* is inductively defined as follows:

1. Any letter a, b, c, d is a nice term.
2. If s_1 and s_2 are nice terms, then so is $[s_1 \circ s_2]$.
3. Nothing else is a nice term.

This definition tells us that something counts as a nice term iff it can be constructed according to the two conditions (1) and (2) in some finite number of steps. In the first step, we construct all nice terms just consisting of letters by themselves, i.e.,

a, b, c, d

In the second step, we apply (2) to the terms we've constructed. We'll get

$[a \circ a], [a \circ b], [b \circ a], \dots, [d \circ d]$

for all combinations of two letters. In the third step, we apply (2) again, to any two nice terms we've constructed so far. We get new nice term such as $[a \circ [a \circ a]]$ —where t is a from step 1 and s is $[a \circ a]$ from step 2—and $[[b \circ c] \circ [d \circ b]]$ constructed out of the two terms $[b \circ c]$ and $[d \circ b]$ from step 2. And so on. Clause (3) rules out that anything not constructed in this way sneaks into the set of nice terms.

Note that we have not yet proved that every sequence of symbols that “feels” nice is nice according to this definition. However, it should be clear that everything we can construct does in fact “feel nice”: brackets are balanced, and \circ connects parts that are themselves nice.

The key feature of inductive definitions is that if you want to prove something about all nice terms, the definition tells you which cases you must consider. For instance, if you are told that t is a nice term, the inductive definition tells you what t can look like: t can be a letter, or it can be $[s_1 \circ s_2]$ for some pair of nice terms s_1 and s_2 . Because of clause (3), those are the only possibilities.

When proving claims about all of an inductively defined set, the strong form of induction becomes particularly important. For instance, suppose we want to prove that for every nice term of length n , the number of $[$ in it is $< n/2$. This can be seen as a claim about all n : for every n , the number of $[$ in any nice term of length n is $< n/2$.

Proposition ind.4. *For any n , the number of $[$ in a nice term of length n is $< n/2$.*

Proof. To prove this result by (strong) induction, we have to show that the following conditional claim is true:

If for every $l < k$, any nice term of length l has $< l/2$ '['s, then any nice term of length k has $< k/2$ '['s.

To show this conditional, assume that its antecedent is true, i.e., assume that for any $l < k$, nice terms of length l contain $< l/2$ '['s. We call this assumption the inductive hypothesis. We want to show the same is true for nice terms of length k .

So suppose t is a nice term of length k . Because nice terms are inductively defined, we have two cases: (1) t is a letter by itself, or (2) t is $[s_1 \circ s_2]$ for some nice terms s_1 and s_2 .

1. t is a letter. Then $k = 1$, and the number of '[' in t is 0. Since $0 < 1/2$, the claim holds.
2. t is $[s_1 \circ s_2]$ for some nice terms s_1 and s_2 . Let's let l_1 be the length of s_1 and l_2 be the length of s_2 . Then the length k of t is $l_1 + l_2 + 3$ (the lengths of s_1 and s_2 plus three symbols '[, \circ,]'). Since $l_1 + l_2 + 3$ is always greater than l_1 , $l_1 < k$. Similarly, $l_2 < k$. That means that the induction hypothesis applies to the terms s_1 and s_2 : the number m_1 of '[' in s_1 is $< l_1/2$, and the number m_2 of '[' in s_2 is $< l_2/2$.

The number of '[' in t is the number of '[' in s_1 , plus the number of '[' in s_2 , plus 1, i.e., it is $m_1 + m_2 + 1$. Since $m_1 < l_1/2$ and $m_2 < l_2/2$ we have:

$$m_1 + m_2 + 1 < \frac{l_1}{2} + \frac{l_2}{2} + 1 = \frac{l_1 + l_2 + 2}{2} < \frac{l_1 + l_2 + 3}{2} = k/2.$$

In each case, we've shown that the number of '[' in t is $< k/2$ (on the basis of the inductive hypothesis). By strong induction, the proposition follows. \square

Problem ind.1. Define the set of supernice terms by

1. Any letter a, b, c, d is a supernice term.
2. If s is a supernice term, then so is $[s]$.
3. If s_1 and s_2 are supernice terms, then so is $[s_1 \circ s_2]$.
4. Nothing else is a supernice term.

Show that the number of '[' in a supernice term t of length n is $\leq n/2 + 1$.

ind.5 Structural Induction

math:ind:sti:
sec

So far we have used induction to establish results about all natural numbers. But a corresponding principle can be used directly to prove results about all **elements** of an inductively defined set. This is often called *structural* induction, because it depends on the structure of the inductively defined objects.

Generally, an inductive definition is given by (a) a list of “initial” **elements** of the set and (b) a list of operations which produce new **elements** of the set from old ones. In the case of nice terms, for instance, the initial objects are the letters. We only have one operation: the operations are

$$o(s_1, s_2) = [s_1 \circ s_2]$$

You can even think of the natural numbers \mathbb{N} themselves as being given by an inductive definition: the initial object is 0, and the operation is the successor function $x + 1$.

In order to prove something about all elements of an inductively defined set, i.e., that every **element** of the set has a property P , we must:

1. Prove that the initial objects have P
2. Prove that for each operation o , if the arguments have P , so does the result.

For instance, in order to prove something about all nice terms, we would prove that it is true about all letters, and that it is true about $[s_1 \circ s_2]$ provided it is true of s_1 and s_2 individually.

Proposition ind.5. *The number of [equals the number of] in any nice term t .*

Proof. We use structural induction. Nice terms are inductively defined, with letters as initial objects and the operation o for constructing new nice terms out of old ones.

1. The claim is true for every letter, since the number of [in a letter by itself is 0 and the number of] in it is also 0.
2. Suppose the number of [in s_1 equals the number of], and the same is true for s_2 . The number of [in $o(s_1, s_2)$, i.e., in $[s_1 \circ s_2]$, is the sum of the number of [in s_1 and s_2 plus one. The number of] in $o(s_1, s_2)$ is the sum of the number of] in s_1 and s_2 plus one. Thus, the number of [in $o(s_1, s_2)$ equals the number of] in $o(s_1, s_2)$. \square

Problem ind.2. Prove by structural induction that no nice term starts with] .

Let’s give another proof by structural induction: a proper initial segment of a string t of symbols is any string s that agrees with t symbol by symbol, read from the left, but t is longer. So, e.g., $[a \circ$ is a proper initial segment of $[a \circ b]$, but neither are $[b \circ$ (they disagree at the second symbol) nor $[a \circ b]$ (they are the same length).

Proposition ind.6. *Every proper initial segment of a nice term t has more ['s than]'s.* math:ind:sti:
prop:initial

Proof. By induction on t :

1. t is a letter by itself: Then t has no proper initial segments.
2. $t = [s_1 \circ s_2]$ for some nice terms s_1 and s_2 . If r is a proper initial segment of t , there are a number of possibilities:
 - a) r is just [: Then r has one more [than it does].
 - b) r is $[r_1$ where r_1 is a proper initial segment of s_1 : Since s_1 is a nice term, by induction hypothesis, r_1 has more [than] and the same is true for $[r_1$.
 - c) r is $[s_1$ or $[s_1 \circ$: By the previous result, the number of [and] in s_1 are equal; so the number of [in $[s_1$ or $[s_1 \circ$ is one more than the number of].
 - d) r is $[s_1 \circ r_2$ where r_2 is a proper initial segment of s_2 : By induction hypothesis, r_2 contains more [than]. By the previous result, the number of [and of] in s_1 are equal. So the number of [in $[s_1 \circ r_2$ is greater than the number of].
 - e) r is $[s_1 \circ s_2$: By the previous result, the number of [and] in s_1 are equal, and the same for s_2 . So there is one more [in $[s_1 \circ s_2$ than there are]. □

ind.6 Relations and Functions

When we have defined a set of objects (such as the natural numbers or the nice terms) inductively, we can also define *relations on* these objects by induction. math:ind:rel:
sec For instance, consider the following idea: a nice term t_1 is a subterm of a nice term t_2 if it occurs as a part of it. Let's use a symbol for it: $t_1 \sqsubseteq t_2$. Every nice term is a subterm of itself, of course: $t \sqsubseteq t$. We can give an inductive definition of this relation as follows:

Definition ind.7. The relation of a nice term t_1 being a subterm of t_2 , $t_1 \sqsubseteq t_2$, is defined by induction on t_2 as follows:

1. If t_2 is a letter, then $t_1 \sqsubseteq t_2$ iff $t_1 = t_2$.
2. If t_2 is $[s_1 \circ s_2]$, then $t_1 \sqsubseteq t_2$ iff $t_1 = t_2$, $t_1 \sqsubseteq s_1$, or $t_1 \sqsubseteq s_2$.

This definition, for instance, will tell us that $a \sqsubseteq [b \circ a]$. For (2) says that $a \sqsubseteq [b \circ a]$ iff $a = [b \circ a]$, or $a \sqsubseteq b$, or $a \sqsubseteq a$. The first two are false: a clearly isn't identical to $[b \circ a]$, and by (1), $a \sqsubseteq b$ iff $a = b$, which is also false. However, also by (1), $a \sqsubseteq a$ iff $a = a$, which is true.

It's important to note that the success of this definition depends on a fact that we haven't proved yet: every nice term t is either a letter by itself, or there are *uniquely determined* nice terms s_1 and s_2 such that $t = [s_1 \circ s_2]$. "Uniquely determined" here means that if $t = [s_1 \circ s_2]$ it isn't *also* $= [r_1 \circ r_2]$ with $s_1 \neq r_1$ or $s_2 \neq r_2$. If this were the case, then clause (2) may come in conflict with itself: reading t_2 as $[s_1 \circ s_2]$ we might get $t_1 \sqsubseteq t_2$, but if we read t_2 as $[r_1 \circ r_2]$ we might get not $t_1 \sqsubseteq t_2$. Before we prove that this can't happen, let's look at an example where it *can* happen.

Definition ind.8. Define *bracketless terms* inductively by

1. Every letter is a bracketless term.
2. If s_1 and s_2 are bracketless terms, then $s_1 \circ s_2$ is a bracketless term.
3. Nothing else is a bracketless term.

Bracketless terms are, e.g., a , $b \circ d$, $b \circ a \circ b$. Now if we defined "subterm" for bracketless terms the way we did above, the second clause would read

If $t_2 = s_1 \circ s_2$, then $t_1 \sqsubseteq t_2$ iff $t_1 = t_2$, $t_1 \sqsubseteq s_1$, or $t_1 \sqsubseteq s_2$.

Now $b \circ a \circ b$ is of the form $s_1 \circ s_2$ with

$$s_1 = b \text{ and } s_2 = a \circ b.$$

It is also of the form $r_1 \circ r_2$ with

$$r_1 = b \circ a \text{ and } r_2 = b.$$

Now is $a \circ b$ a subterm of $b \circ a \circ b$? The answer is yes if we go by the first reading, and no if we go by the second.

The property that the way a nice term is built up from other nice terms is unique is called *unique readability*. Since inductive definitions of relations for such inductively defined objects are important, we have to prove that it holds.

Proposition ind.9. *Suppose t is a nice term. Then either t is a letter by itself, or there are uniquely determined nice terms s_1, s_2 such that $t = [s_1 \circ s_2]$.*

Proof. If t is a letter by itself, the condition is satisfied. So assume t isn't a letter by itself. We can tell from the inductive definition that then t must be of the form $[s_1 \circ s_2]$ for some nice terms s_1 and s_2 . It remains to show that these are uniquely determined, i.e., if $t = [r_1 \circ r_2]$, then $s_1 = r_1$ and $s_2 = r_2$.

So suppose $t = [s_1 \circ s_2]$ and also $t = [r_1 \circ r_2]$ for nice terms s_1, s_2, r_1, r_2 . We have to show that $s_1 = r_1$ and $s_2 = r_2$. First, s_1 and r_1 must be identical, for otherwise one is a proper initial segment of the other. But by **Proposition ind.6**, that is impossible if s_1 and r_1 are both nice terms. But if $s_1 = r_1$, then clearly also $s_2 = r_2$. \square

We can also define functions inductively: e.g., we can define the function f that maps any nice term to the maximum depth of nested $[\dots]$ in it as follows:

Definition ind.10. The *depth* of a nice term, $f(t)$, is defined inductively as follows: mth:ind:rel:
defn:depth

$$f(t) = \begin{cases} 0 & \text{if } t \text{ is a letter} \\ \max(f(s_1), f(s_2)) + 1 & \text{if } t = [s_1 \circ s_2]. \end{cases}$$

For instance

$$\begin{aligned} f([a \circ b]) &= \max(f(a), f(b)) + 1 = \\ &= \max(0, 0) + 1 = 1, \text{ and} \\ f([[a \circ b] \circ c]) &= \max(f([a \circ b]), f(c)) + 1 = \\ &= \max(1, 0) + 1 = 2. \end{aligned}$$

Here, of course, we assume that s_1 and s_2 are nice terms, and make use of the fact that every nice term is either a letter or of the form $[s_1 \circ s_2]$. It is again important that it can be of this form in only one way. To see why, consider again the bracketless terms we defined earlier. The corresponding “definition” would be:

$$g(t) = \begin{cases} 0 & \text{if } t \text{ is a letter} \\ \max(g(s_1), g(s_2)) + 1 & \text{if } t = s_1 \circ s_2. \end{cases}$$

Now consider the bracketless term $a \circ b \circ c \circ d$. It can be read in more than one way, e.g., as $s_1 \circ s_2$ with

$$s_1 = a \text{ and } s_2 = b \circ c \circ d,$$

or as $r_1 \circ r_2$ with

$$r_1 = a \circ b \text{ and } r_2 = c \circ d.$$

Calculating g according to the first way of reading it would give

$$\begin{aligned} g(s_1 \circ s_2) &= \max(g(a), g(b \circ c \circ d)) + 1 = \\ &= \max(0, 2) + 1 = 3 \end{aligned}$$

while according to the other reading we get

$$\begin{aligned} g(r_1 \circ r_2) &= \max(g(a \circ b), g(c \circ d)) + 1 = \\ &= \max(1, 1) + 1 = 2 \end{aligned}$$

But a function must always yield a unique value; so our “definition” of g doesn’t define a function at all.

Problem ind.3. Give an inductive definition of the function l , where $l(t)$ is the number of symbols in the nice term t .

Problem ind.4. Prove by structural induction on nice terms t that $f(t) < l(t)$ (where $l(t)$ is the number of symbols in t and $f(t)$ is the depth of t as defined in [Definition ind.10](#)).

Photo Credits

Bibliography