Gold's Theorems and Locking Sequences

Nina Gierasimczuk   Dick de Jongh

February 8th, 2013

If we consider a class $\mathcal{L}$ of languages here, we assume it to be numbered, i.e. each language has a number. Later we will be more precise about the numbering.

Theorem 1 (Gold 1967) *The class of all finite languages is identifiable*[1].

Proof Define $M(\sigma) = content(\sigma)$ (i.e., the set of elements contained in σ; of course, we should properly say, the number of a language $L \in \mathcal{L}$, such that $L = content(\sigma)$.) Let L be finite and t a text for L. We need to show that M converges on t to a number for L. For some large enough n, all elements of L will have occurred in $t[n]$, i.e. $content(t[n]) = L$, and it will remain that way, $content(t[m]) = L$ for all $m > n$. So, M converges to a number for L. □

A *cofinite* subset of the natural numbers (or more sloppily cofinite set) is a set that contains all elements of $\mathbb{N}$ except finitely many. Similarly, we can talk about co-singleton sets or co-doubleton sets.

Example 2 *The class of all co-singleton sets* $\mathcal{L} = \{L_i \mid i \in \mathbb{N}\}$ *with* $L_i = \mathbb{N} - \{i\}$ *is identifiable.*

Proof Take $M(\sigma) = \mathbb{N} - \mu n(n \notin content(\sigma))$. Let t be a text for $\mathbb{N} - \{k\}$. For a certain m, $\{0, \ldots, k-1\} \subseteq content(t[m])$. For all $m' \geq m$, the minimal number not in $content(t[m'])$ will be k, so $M(t[m']) = k$ and M converges to $\mathbb{N} - \{k\}$'s number. □

Theorem 3 (Gold 1967) *Let $\mathcal{L}$ be a family of languages that contains all finite ones and at least one infinite one. Then $\mathcal{L}$ is not identifiable.*

Proof Let $\mathcal{L}$ contain the infinite set A and all finite sets. We will, for any M that is presumed to identify $\mathcal{L}$ construct a text t for A on which M fails to converge. Let $x_1, x_2, x_3, \ldots$ enumerate A. The text t is constructed in stages. In stage n the initial segment σ_{n-1} of t that has been constructed will be extended to an initial segment σ_n in such a way that $content(\sigma_n) = \{x_1, x_2, \ldots, x_n\}$, i.e. exactly $x_1, x_2, \ldots, x_n$ are used in σ_n.

[1] Unless specified otherwise, by 'identifiable' we mean 'identifiable in the limit'.

- *Stage* 1. Consider the text $x_1, x_1, x_1, \ldots$. This is a text for the finite language $\{x_1\}$ supposedly identified by M. So, for a certain initial segment σ_1 of this text, $M(\sigma_1) = \{x_1\}$. This will be the initial segment σ_1 of t.

- *Stage* $n+1$. Let the part of t constructed in stage n be σ_n with $content(\sigma_n)= \{x_1, x_2, \ldots, x_n\}$. Consider the text $\sigma_n{}^\wedge x_{n+1}, x_{n+1}, x_{n+1}, \ldots$. This is a text for the language $\{x_0, x_1, \ldots, x_n, x_{n+1}\}$. So, M will, after σ_n, sooner or later give $\{x_1, x_2, \ldots, x_{n+1}\}$ as a value. If indeed $M(\sigma_n{}^\wedge\langle x_{n+1}, \ldots x_{n+1}\rangle) = \{x_1, x_2, \ldots, x_{n+1}\}$, we take that string $\sigma_n{}^\wedge\langle x_{n+1}, x_{n+1}, \ldots x_{n+1}\rangle$ to be σ_{n+1}. It is obvious that t is a text for A and M does not converge on t.

$\square$

The second Gold theorem can be generalized easily. We leave this as an exercise. At the heart of this proof an important concept is hidden, the concept of locking sequence.

Definition 1 (Blum and Blum 1975) *A sequence σ is a* locking sequence *for learner M and language L if $content(\sigma) \subseteq L$ and, for each τ with $content(\tau) \subseteq L$, $M(\sigma^\wedge\tau) = M(\sigma)$.*

Theorem 4 (Locking sequence theorem) *If M identifies L, then there exists a locking sequence σ for L such that $M(\sigma)$ is a number for L.*

Before we prove this theorem let us show how to prove Gold's second theorem quickly by applying the locking sequence theorem.

Proof Let σ be a locking sequence for M and A such that $M(\sigma)$ is a number for A, and let x be the first element of σ. Consider the text $t = \sigma^\wedge x, x, x, \ldots$. This is a text for the finite set *content*(σ). For any sequence of x's, $M(\sigma^\wedge\langle x, \ldots, x\rangle) = M(\sigma)$, so M converges on t to the number of A, and does not identify *content*(σ), contradicting the fact that M identifies all finite sets. $\square$

Proof of Theorem 4. Assume M identifies L without there being a locking sequence for M and L on which M gives a number for L.

Let us first assume that a locking sequence σ does exist, but that $M(\sigma)$ is not a number for L. This is clearly absurd, since the M would keep giving a number that is not a number for L on any text for L that starts with σ.

Now it will be sufficient to show that a contradiction follows from the assumption that there exists no locking sequence at all. We construct in stages a text t for L on which M does not converge. Let $x_1, x_2, x_3, \ldots$ enumerate L.

- *Stage* 1. The string $\langle x_1\rangle$ is not a locking sequence, so for some τ over L, $M(\langle x_1\rangle^\wedge\tau) \neq M(\langle x_1\rangle)$. Take $\langle x_1\rangle^\wedge\tau$ as the initial segment σ_1 of t.

- *Stage* $n+1$. Assume the initial segment σ_n of t has been constructed in stage n. By assumption, the sequence $\sigma_n{}^\wedge\langle x_{n+1}\rangle$ is not a locking sequence, so there is a sequence τ over L such that $M(\sigma_n{}^\wedge\langle x_{n+1}\rangle^\wedge\tau) \neq M(\sigma_n{}^\wedge\langle x_{n+1}\rangle)$. Take $\sigma_{n+1} = \sigma_n{}^\wedge\langle x_{n+1}\rangle^\wedge\tau$.

Because each x_i occurs in t, t is a text for L. But learner M keeps changing value on t, it does not converge. $\square$

Example 5 *The class $\mathcal{L}$ containing the co-singleton sets, and in addition the set $L_0 = \mathbb{N}$ is not identifiable.*

Proof Assume M identifies $\mathcal{L}$. Let σ be a locking sequence for M and L_0. Let k be the minimal natural number not in *content*(σ). Consider the text $t = \sigma^\wedge\langle 0, 1 \ldots, k-1, k+1, k+2, k+3, \ldots\rangle$. On every initial segment of t from σ onwards, M will give the value 0. But t is a text for an $L \in \mathcal{L}$, namely $L = \mathbb{N} - \{k\}$. Therefore M does not identify $\mathbb{N} - \{k\}$ and hence neither $\mathcal{L}$. $\square$

The locking sequence theorem can be used to obtain a characterization of the identifiable collections of languages.

Definition 2 (Angluin 1980) *Let L be a member of the collection of languages $\mathcal{L}$. A finite subset D of L is a* telltale subset *of L w.r.t. $\mathcal{L}$ if it has the property that*

$$\forall L' \in \mathcal{L}\,(D \subseteq L' \Rightarrow L' \not\subset L)$$

or equivalently, $\forall L' \in \mathcal{L}\,(D \subseteq L' \subseteq L \Rightarrow L' = L)$.

Theorem 6 (Angluin 1980) *A class of languages $\mathcal{L}$ is identifiable iff each $L \in \mathcal{L}$ has a telltale subset D_L.*

Proof
[$\Rightarrow$] Let M identify $\mathcal{L}$. Consider a locking sequence σ for M and L, and take $D_L = \mathit{content}(\sigma)$. We will show that D_L is a telltale subset of L. Assume that it is not, i.e., $D_L \subseteq L' \subset L$ for $L' \in \mathcal{L}$. It suffices to get a contradiction.

Assume $x_1, x_2, x_3, \ldots$ is an enumeration of L'. Consider the text $t = \sigma^\wedge\langle x_1 x_2, x_3, \ldots\rangle$. This is a text for L'. Since t starts with a locking sequence for L and contains only elements from L, M will converge on t to an number for L, which contradicts the fact that it is supposed to identify L' as well.

[$\Leftarrow$] Assume each $L \in \mathcal{L}$ has a telltale subset D_L. Define M in the following way:

$M(\sigma) = \mu e'(e'$ is a number for some $L' \in \mathcal{L}$ such that
$D_{L'} \subseteq \mathit{content}(\sigma) \subseteq L')$ if such e' exists, and 0 otherwise.

Assume t is a text for L and e is the least number for L. It is sufficient to show that, for k large enough, $M(t[k]) = e$. Fix n large enough so that $D_L \subseteq \mathit{content}(t[n])$. As t is a text for L, L and e now satisfy
e is the minimal number for L and $D_L \subseteq \mathit{content}(t[n]) \subseteq L$.
Nevertheless, we cannot conclude that $M(t[n]) = e$, because there may be (finitely many) other languages $L_1, \ldots, L_m$ with indices $e_1, \ldots, e_m < e$ (and therefore different from L) that satisfy the same condition $(1 \leq i \leq m)$:

e_i is the minimal number for L_i and $D_{L_i} \subseteq content(t[n]) \subseteq L_i$.

Then $M(t[n])$ would be the smallest of these pretenders e_i. Take any such language L_i. By the telltale condition, since we now also have $D_{L_i} \subseteq content(t[n]) \subseteq L$, and in particular therefore $D_{L_i} \subseteq L$, there will be an $x_i \in L$ not in L_i. As t is a text for L there will be a k_i such that $x_i \in t[k_i]$. We now take k to be the maximum of all the k_i and i, and then only L will satisfy $D_L \subseteq content(t[k]) \subseteq L$, and this will remain so for numbers $> k$. So, all the L_i have been eliminated and M will keep producing e as its value. $\square$

References

Angluin, D. (1980). Inductive inference of formal languages from positive data. *Information and Control*, 45(2):117–135.

Blum, L. and Blum, M. (1975). Toward a mathematical theory of inductive inference. *Information and Control*, 28:125–155.

Gold, E. M. (1967). Language identification in the limit. *Information and Control*, 10:447–474.