

A COGNITIVE MODEL OF THE LEARNING CURVE

Guido Fioretti
CEEL, Trento, Italy
ICER, Torino, Italy

February 7, 2001

Abstract

This article provides a cognitive foundation of the parameters that regulate a model of the learning curve. Organizational learning and its actual occurrence are linked to the number of available categories and to the amount of information to be processed.

Guido Fioretti
via San Massimo 8
10123 Torino
Italy
Fax: ++390744425713
E-mail: fellow1@icer.it
guido.fioretti@lycos.com

1 Introduction

The learning curve denotes a rather robust empirical regularity, namely that production time decreases with cumulative production at a uniform rate [11]. In other words, the more units of a good have been produced, the less it takes to produce an additional unit.

Empirical learning curves are roughly described by the following equation:

$$t_n = t_1 N^{-\alpha} \tag{1}$$

where t_n is the time required to produce the n -th unit, t_1 is the time required to produce the first unit, $N = 1 + 2 + \dots + n$ is cumulative production and α is a parameter that is specific to the manufacturing process being observed.

The most important finding of empirical research on learning curves is that production time reduction takes place to a much larger extent in assembling operations than in machining operations [2, 3]. This evidence suggests that the learning curve is not concerned with individual learning, but rather with organisational learning.

Evidently, the learning curve does not arise from workers learning to use a new machine or employees learning to deal with a new boss. Rather, it arises from individuals learning to coordinate their actions. Possibly, it is not a chance that the learning curve was discovered in the aircraft industry, where specialists from different branches must learn how to fit together many different components.

Many other empirical facts point to the organizational origin of the learning

curve. One such fact is that if the management sets a ceiling to possible improvements of production time, learning does not occur beyond that ceiling [4]. Other interesting facts are that learning rates may differ across plants of the same firm that produce the same good with similar equipment within the same country, and that production interruptions such as strikes cause production time to increase [1].

The above fact about managers setting a ceiling to learning suggests that the learning curve has a lot to do with workers motivation and with operating in an environment that provides stimuli to learn. However, the story about learning rates differing across plants suggests that the learning curve does not derive from individuals learning a given corporate culture. Moreover, the empirical fact about strikes suggests that the learning curve originates from tacit coordination schemes that may be difficult to reconstruct, rather than codified behavior rules that can be easily re-learned.

Muth [8] provided a model of the learning curve based on random search among a fixed population of possible technologies. However, a technology-based model is not able to explain why learning occurs to a much larger extent in assembling operations than in machining operations, as well as the many other empirical facts mentioned above.

Alternatively, Huberman derived the learning curve from the properties of a random graph that represents the structure of communications within a firm [9, 7]. In this model, nodes represent organizational units and edges represent communication between units. Organizational units may be human beings, machines, or compounds of both. Communication between units may take place in connection

with physical transportation of goods, or it may be a pure information flow.

In a random graph, nodes establish links to one another with a certain probability. According to the above metaphor, this is the probability that organizational units establish communications with one another.

Huberman's model assumes that communication within the firm must flow from a source node to a sink node, corresponding to the flow of production. Clearly, this model neglects power hierarchies within firms. However, assembling and coordination operations are likely to take place between organizational units placed at similar hierarchical levels, so for our purposes hierarchy can be ignored.

Since production corresponds to a communication flow from source node to sink node, production time can be represented by the number of steps that it takes to cross the graph. Learning takes place when units understand how to communicate with one another, that is when the nodes of the graph establish links to one another. Thus, knowledge stemming from cumulative production is represented by p , the probability to establish a link between two nodes.

Huberman's model yields a relationship between cumulative production and production time that resembles 1. Moreover, it introduces a parameter r that explains why the learning curve occasionally fails.

Parameter r is the probability of eliminating unproductive connections when looking for a path from the source to the sink of the graph. It represents the effectiveness of the search for better coordination. Basically, it is an index of the quality of decision-making.

Interestingly, at high values of r production time decreases exponentially with p , as it is empirically observed when learning curves work. However, the lower r , the weaker this relationship. At $r = 0.5$ the relationship is reverted, with production time increasing with cumulative output.

Interpretation is straightforward. The learning curve arises from communication between organizational units, represented by p . That's why it is more pronounced in assembling operations than in machining operations. However, setting a ceiling to production time slackens the search for better coordination, reduces r , and causes the learning curve to fail. Obviously, firm-specific values of r can easily explain why the rate of learning eventually differs across plants of the same firm that produce the same good with similar equipment within the same country. Finally, after strikes or other production interruptions tacit knowledge concerning coordination is difficult to reconstruct, whereas explicit knowledge concerning how to operate specific machines can be easily acquired.

The above discussion highlighted that Huberman's model is able to shed light on the (mis)functioning of the learning curve. However, this model rests on parameters that are not linked to the process of learning. Huberman's model gives no hint as to which job could exhibit a high p or a low r , and why. A cognitive foundation of Huberman's model is needed, and the rest of this paper sets out to provide it.

2 The Market for Ideas


Learning means recognizing that certain situations have common features, and developing an appropriate behavior for each class of situations that a decision-maker expects to face. Learning involves the cognitive process of classifying information into a manageable number of mental categories, and the evolutionary process of developing decision rules that yield favourable results.

One such model is John Holland's *Classifier Systems* [5, 6]. Let us adapt it to a random graph representing the structure of communications within a firm.

Let us assume that the organizational units of our firm are endowed with categories which they use to classify information, plus a rule that specifies on which occasion a particular piece of information should be produced. Since organizational units generally represent compounds of men and machines, categories represent the situations that decision-makers endowed with particular machines are able to distinguish, while the information that they produce represents the decisions they make.

Let us represent both categories and information by means of strings of L characters that can be zeros, ones, or "don't care" characters #. Thus, $K = 3^L$ different strings can be produced.

Categories must entail at least one #-character in order to be effective. In fact, categories are strings that match all information strings that have zeros and ones in the same positions where they have zeros and ones; on the contrary, it does not matter which characters information strings have, in the positions where category



a category string information strings matched by that category string

Figure 1: A category string, and the four information strings that it is able to match. This category string does not distinguish these information strings from one another. Thus, four information strings are classified as one single piece of information.

strings have a #. In this way, a category string is a container that classifies all information strings that have zeros and ones in its same positions. For instance, the category string depicted in the left half of figure 1 is able to classify the four information strings depicted in the right half of the picture.

Information strings may have #-characters as well. However, these characters have a different meaning as in category strings.

Let us consider a generic organizational unit j . If an information string produced by unit j has #-characters in the same positions where a category string owned by j has #-characters, then unit j simply carries on to other units the symbol that was matched by its category. That is, if the category string had a # in a position that was matched by a 0 (1), then the information string produced by j

has also a 0 (1) in that position. In this way, information clusters can propagate.

A classifier system works like a market. Category strings bid "prices" in order to classify information strings, and they pay if they succeed to match information strings. Category strings pay with a conventional currency called *strength*. Strength represents the value of information for a firm. If a firm's organizational units wants to ignore certain information, it simply doesn't bid any "strength" for it.

Strength flows from units that receive information to units that issue it. Organizational units, in their turn, distribute their endowment of "strength" among their categories.

In our case, it is a "market" for ideas about coordinating activities across organizational units. It is supposed to take place among organizational units placed at similar hierarchical levels, that must coordinate in order to carry out the kind of assembling operations where learning curves arise. Each organizational unit tries to understand what other units need, and presents proposals concerning the way of carrying out its own particular job. Category strings represent the ability of organizational units to understand the proposals of other units. Information strings represent the proposals they make.

By passing across organizational units, strength creates preferential paths for information flows. Moreover, if strength is passed on along a circular path, the same sequence of actions is repeated over and over. This is what happens when organizational units learn to coordinate.

Strength cannot be the only magnitude to decide which category string will

classify a certain information string. In fact, categories entailing a large number of #-characters are inherently better at classifying information — in the limit, a category made only by #-characters would be able to classify any information. However, a general-purpose category would represent a decision-maker who is not able to distinguish pieces of information from one another.

Consequently, classifier systems employ two magnitudes in order to decide which category string classifies which information string. The second magnitude is the *specificity* of a string, defined as the number of its non-#-characters.

Let us summarize the above considerations as follows. Let indices $k, k' = 1, 2, \dots, K$ denote category strings and information strings, respectively. Let $p_{kk'}$ denote the probability that a category string of type k classifies an information string of type k' . Let $t_{kk'}$ denote the strength that category string k passes on to information string k' . Let s_k be the specificity of category string k and let $s_{k'}$ be the specificity of information string k' , where $s_{k'} \geq s_k$. Then, $p_{kk'}$ must be such that $\partial p_{kk'} / \partial t_{kk'} > 0$ and $\partial p_{kk'} / \partial (s_{k'} - s_k) < 0$.

Following the multinomial logit model, let us choose the following expression:

$$p_{kk'} = \frac{e^{\beta \frac{t_{kk'}}{s_{k'} - s_k}}}{\sum_{k'} e^{\beta \frac{t_{kk'}}{s_{k'} - s_k}}} \quad (2)$$

Formula 2 concerns the probability of establishing a connection between one category string of a certain type and one information string of a certain type. The probability of establishing a connection between two organizational units, in its turn, depends on the type of strings the two units are endowed with, as well as on

their strength endowments.

Taking a small step away from Holland's classifier systems, let us assume that the management endows organizational units with additional strength if they happen to fall short of it. In this way, we are assuming that organizational units can always bid the amount of strength they desire.

Clearly, strengths no longer obey a conservation law. However, new amounts of strength are not created arbitrarily. They represent occasional interventions by the management in order to avoid that some units lag behind. Remember that we are describing a firm, not a market.

With this proviso to the credit assignment algorithm, the probability of establishing a connection between a category string of type k and an information string of type k' equals the probability of establishing a connection between two organizational units endowed with k and k' , respectively. That is, for any pair of organizational units i and j we have

$$P_{i(k)j(k')} \equiv P_{kk'} \tag{3}$$

where $i(k)$ means that unit i is endowed with category string k and unit j is endowed with information string k' .

3 Equilibrium Knowledge

In the previous section, learning was understood as a process of information classification. Information strings and category strings had been assumed to be given, so learning consisted of choosing the best categories for any given information.

This picture is correct, but incomplete. It is the picture of a firm whose organizational units are concerned with accomplishing given tasks using given means, without ever influencing or modifying these tasks and these means. It is a bureaucratic picture of a firm, describing the execution of orders along a hierarchy.

However, the learning curve typically arises from assembling operations that involve units adapting their expertise in order to coordinate with one another. Learning takes the form of mutual understanding of one other's problems, rather than implementation of the management's directives. In this context, tasks and means of every unit are flexible, and their choice is up to the unit itself.

In our model, flexibility of tasks and means corresponds to allowing category strings and information strings to evolve with time. This is absolutely normal for classifier systems, where new strings are either generated by mutation of existing strings or by random recombination of parts of them (*cross-over*).

Generation of new strings is of the utmost importance for classifier systems, since it ensures that optimal strings can be found. Interestingly, both mutation and cross-over bear similarities to firm practices. On the one hand, mutation reflects invention and introduction of novel practices. On the other hand, cross-over resembles communication along standardized formats and innovation through re-

combination of working solutions [10].

Let us consider the equilibrium state of the evolutionary dynamics that creates new strings. At equilibrium, optimal strings have been found. In this state, evolution has selected category strings and information strings such that $t_{kk'}$ and $(s_{k'} - s_k)$ take the same values for all k and k' . Thus, probabilities $p_{kk'}$ take the same values for all k, k' .

Let H denote the number of different category strings that are employed by the firm, and let H' denote the number of different information strings that are produced by the firm. Since $K = 3^L$ is the number of different strings that can be generated out of L symbols, it is $H \leq K$ and $H' \leq K$. Furthermore, since the purpose of a category string is that of classifying many information strings, it must be $H < H'$.

Since $t_{kk'}$ and $(s_{k'} - s_k)$ take the same values for all k and k' , and since index k' extends to H' , from 2 we obtain $p_{kk'}^* = 1/H'$, where the asterisk is there to remind that this is an equilibrium value.

This is the equilibrium probability that a category string of type k matches an information string of type k' , for $\forall k, k'$. Because of 3, this is also the probability of establishing a link between any node i that owns a category string of type k and any node j that issues an information string of type k' , for $\forall k, k'$.

Thus, the probability that a unit i establishes a link with whatever unit j , independently of the category owned by i , is the sum of $1/H'$ over the H category types. This yields H/H' , which we can take as equivalent to the p parameter of

Huberman's model:

$$p = \frac{H}{H'} \quad (4)$$

The r parameter, in its turn, represents the effectiveness of the search for a path from the source node to the sink node. Since a simple measure of effectiveness is that trials are not duplicated, we can identify r with the probability that any two nodes are *not* linked by multiple paths.

Given 4, two nodes are linked by two paths with probability $(H/H')^2$, by three paths with probability $(H/H')^3$, and so on. Thus, the probability that two nodes are linked by at least two paths is the sum of the series $(H/H')^2 \sum_{i=0}^{\infty} (H/H')^i$. Since it is $\sum_{i=0}^{\infty} x^i = 1/(1-x)$, the sum of that series is $H^2/H'(H' - H)$. This can be divided by H in order to yield numbers that are less than one, obtaining $H/H'(H' - H)$.

Consequently, the r parameter is given by $1 - H/H'(H' - H)$, which is equivalent to:

$$r = \frac{H'^2 - HH' - H}{H'^2 - HH'} \quad (5)$$

Equations 4 and 5 above link the parameters that regulate the shape of the learning curve to cognitive features of the firm, namely the number of different category strings and the number of different information strings. The next section will explain the meaning of formulas 4 and 5 by means of a numerical example.




Figure 2: Parameters p and r for $H = 1, 2, \dots, 10$, with $H' = H + 1$.

4 A Numerical Example

Let us assume $H' = H + 1$, which is the minimum difference between the number of different category strings and the number of different information strings. Since the difference between H and H' is so small, let us consider small values of H as well. Figure 2 depicts parameters p and r for $H \in [1, 10]$ and $H' \in [2, 11]$.

For any endowment of category strings and information strings, figure 2 provides equilibrium connection probability p . This is the probability of establishing a link between any two organizational units that is attained when the firm optimizes utilization of cognitive resources. Consequently, p can be considered as a kind of learning potential.

Thus, figure 2 is telling us that the greater the number of different category strings and the greater the number of information strings, the higher the learning

potential of a firm. The more cognitive resources are available, the more can be learned.

However, learning potential is wasted if the search for a better organization of jobs is ineffective. Parameter r measures the effectiveness of the search for better coordination. Figure 2 shows that the greater the number of different category strings and the greater the number of information strings, the lower the effectiveness of the search for jobs coordination.

Figure 2 illustrates a trade-off between what can be learned, and what is actually learned. The more is there to learn, the more difficult it is that a firm actually learns it. With $H' = H + 1$ parameter r is all the time below the 0.5 threshold, where — according to Huberman's model — the learning curve fails to materialize.

Clearly, the above results are specific to our choice of $H' = H + 1$. A different ratio between the number of different category strings and the number of different information strings is likely to yield different results.

Possibly, a cognitive system that would be able to deal with the same number of different information strings with a smaller number of category strings, would be more efficient. If we can simplify a lot of information by means a small number of categories, that's when we easily recognize what kind of problem we are facing and what kind of solutions can be applied.

Figures 3 and 4 show the effect of switching from $H' = H + 1$ to $H' = H + 2$ on p and r , respectively. Even if we decreased the number of different categories by only one unit, we can observe dramatic differences. Learning potential p de-




Figure 3: Parameter p when $H = H' - 1$ (thin line) and when $H = H' - 2$ (thick line).

creases, but the higher the number of different strings, the less it decreases. On the contrary, learning effectiveness represented by r increases, and the higher the number of different strings, the more pronounced this effect.

Figures 3 and 4 suggest that by limiting the number of different categories a large gain of learning effectiveness can be obtained at the expense of a little loss of learning potential. A limited number of well-designed categories, able to discriminate relevant features but ultimately compressing a lot of information into manageable limits, is likely to do better than any attempt to use all information available to a decision-maker. Bounded rationality and thumb rules do not merely arise out of limitations of our capabilities, but constitute a basis of efficient decision-making as well.




Figure 4: Parameter r when $H = H' - 1$ (thin line) and when $H = H' - 2$ (thick line).

5 Conclusions

Throughout the previous sections, this article provided a foundation of a model of the learning curve in terms of cognitive features of the units involved in organizational learning within a firm. These features are the number of different category strings and the number of different information strings.

Two results have been established. The first one is that by increasing the number of different category strings and information strings, learning potential increases but learning effectiveness decreases. The more complex a problem, the more can be learned, but the more difficult it is to learn it. The second one is that by using a few categories to classify a lot of information, it is possible to increase learning effectiveness at the expense of a little loss of learning potential. If one is able to solve complex problems by means of simple schemes and thumb rules,

that's the best.

Qualitatively, the above results are obvious. However, these results have been formalized, and formalization may prelude to quantitative application. Information strings represent the actions that the organizational units of a firm can undertake, category strings represent the interpretations that these units give of the information they receive. Behavioral routines and technological possibilities can be codified into strings of zeros, ones, and "don't care" characters. Eventually, potential gains from organizational learning could be estimated, and the reliability of a forecast based on the learning curve could be assessed.

Nonetheless, one should never forget that the above analysis is affected by serious limitations. Semantics, i.e. the connection between a string of symbols and a real action that has a precise meaning for the firm that is performing it, lies out of the scope of this model. It is up to the modeller to attach a meaning to the strings that are in the model. In other words, this model can describe organizational learning out of recombination of existing building blocks, but cannot describe the invention of new conceptual blocks.

References

- [1] Linda Argote and Dennis Epple. Learning curves in manufacturing. *Science*, 247:920–924, 1990.
- [2] Werner Z. Hirsch. Manufacturing progress functions. *The Review of Economics and Statistics*, 34:143–155, 1952.
- [3] Werner Z. Hirsch. Firm progress ratios. *Econometrica*, 24:136–143, 1956.

- [4] Winfred B. Hirschmann. Profit from the learning curve. *Harvard Business Review*, 42:125–139, 1964.
- [5] John H. Holland. *Adaptation in Natural and Artificial Systems*. The University of Michigan Press, Ann Arbor, 1975.
- [6] John H. Holland. Escaping brittleness: The possibilities of general-purpose learning algorithms applied to parallel rule-based systems. In *Machine Learning: An Artificial Intelligence Approach*, volume 2, Los Altos, 1986. Morgan Kaufmann Publishers.
- [7] Bernardo A. Huberman. The dynamics of organizational learning. Technical report, Xerox Palo Alto Research Center, 1997. Available at <http://www.parc.xerox.com/istl/groups/iea/abstracts/MultiagentDynamics/orglearn.html>.
- [8] John F. Muth. Search theory and the manufacturing progress function. *Management Science*, 32:948–962, 1986.
- [9] Jeff Shrager, Tad Hogg, and Bernardo A. Huberman. A graph-dynamic model of the power law of practice and the problem-solving fan-effect. *Science*, 242:414–416, 1988.
- [10] Durward K. Sobek, II, Jeffrey K. Liker, and Allen C. Ward. Another look at how toyota integrates product development. *Harvard Business Review*, 76:36–49, 1998.
- [11] Louis E. Yelle. The learning curve: Historical review and comprehensive survey. *Decision Sciences*, 10:302–328, 1979.