Chapter 17

Learning and forgetting communicative messages

Jaap Murre & Antonio Chessa

Introduction

"Contrary to the existing law, from now on cyclists coming from the right on an intersection have right of way over cars." Communicating such a message stands at the beginning of a long chain of psychological and neurobiological events that ends with behaviour. After being exposed to the message, we must first encode it. The type and depth of encoding depend on the learning circumstances and on our prior knowledge. After encoding, a representation of the message is stored in the brain's memory processes and structures. These processes are volatile, and a memory may decay spontaneously. Learning of similar messages may also interfere with continued storage. At any time, retrieval of the message may be attempted, possibly cued by a specific situation. We approach an intersection and notice a boy on a cycle coming from the right. The message about the new law pops into our consciousness. We brake and a fatal accident is averted. After many such successful retrievals and actions, a habit forms. We then brake unconsciously without first having to retrieve the message. Only now has the message reached its final goal: it has become a habit. On the way to its final goal, we can distinguish three principal stages of memory, namely encoding, storage and retrieval (see Figure 1). We will discuss these three stages below.

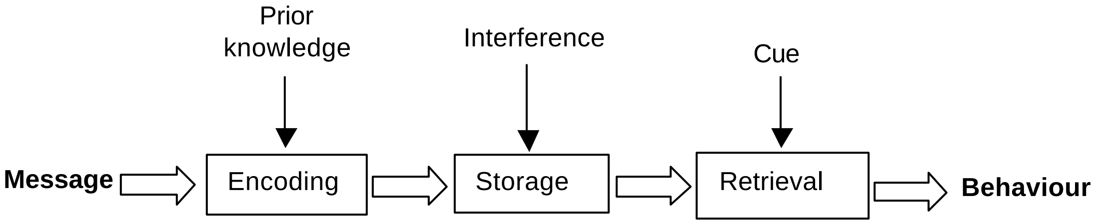

Figure 1: The three principal stages of memory: encoding, storage and retrieval.

The empirical study of learning and forgetting starts with the work of the German scholar Hermann Ebbinghaus who in 1885 taught himself lists of carefully constructed nonsense syllables and then measured how long it took to relearn them after a period of time. The percentage of time saved on relearning is called *savings*. The more time that has elapsed, the lower the savings, which can be visualized in a forgetting curve (see Figure 2).

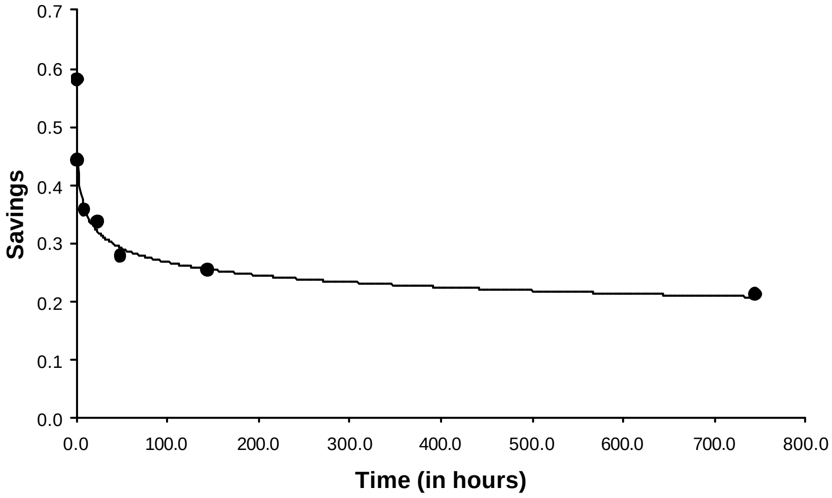

Figure 2: The forgetting curve by Ebbinghaus (1885). The dots indicate the original data; the solid curve is a trend line fitted to the data (a power function).

Over the past century, psychologists and neurobiologists have uncovered a wealth of facts and we now understand many of the brain's memory processes. This chapter will introduce the principal findings in this area and touch upon recent advances in theory development. As we will see, many of these results are relevant for the effective and enduring communication of messages.

Encoding and learning

Factors affecting encoding

A message reaches us in a physical appearance, which is processed through one or more modalities. For example, we see and hear an item about a new law on the evening news. What we remember of such a message depends on the mental processing carried out when we perceive it. Craik and Lockhart discovered in 1972 that when subjects are instructed to pay attention to the *physical* form of a message they will remember the message for a shorter time compared with others who are instructed to attend to the *meaning* of the message. So, if we are merely looking for a specific sequence of letters such as 'ghe' (a low-level task), in a list of words, we are less likely to retrieve the list than when we are assessing each word on the list for pleasantness (a high-level task). Craik and Lockhart argued, on the basis of experiments like this, that higher levels of processing lead to deeper and longer-lasting encoding (also see the ELM model, described in Chapter 16 in this book by Petty).

Many factors contribute to the 'memorability' of a message. Novelty, defined as the degree to which we do *not* expect something, typically leads to longer-lasting encoding. So do items which stand out from their surroundings. One could argue that the brain should encode only that which it does not expect and which it cannot infer.

In the past 50 years, neuroscientists have discovered many details of the brain's encoding process. For example, certain neuromodulatory substances must be secreted by the brain in order for a memory to be laid down. Blocking of their working will prevent memory formation: subjects may not remember anything from the experimental episode. There is evidence that the brain secretes more neuromodulators in response to unexpected events. Other global physiological states of wakefulness and arousal also influence their availability and thus the likelihood of reliable encoding of a message. Raising arousal levels with loud noise or perceived danger will increase the chances of remembering a message presented simultaneously, even if the source of the arousal is unrelated to the message (also see Chapter 9 in this book by De Vries & Leegwater).

Implicit versus explicit memory

One could argue that most encoding occurs unconsciously. Merely by being exposed to the world around us, our behaviour is influenced in subtle ways, without us being aware of it. An unconscious influence on our behaviour is called *implicit memory* (Schacter, 1987). Explicit memory, its counterpart, refers to those memories that we can consciously recall, for example, when told about the date, place, and when given one or two additional cues: When we had dinner in Amsterdam last summer and the waiter dropped a whole serving tray in the canal.

Typical tasks used to study implicit memory are word stem completion and category generation. With stem completion, a word stem such as WIN___ is presented with the instruction to complete it to form any word that comes to mind. There is evidence of implicit memory when there is an increased tendency to complete the stem to a word that has been presented earlier, for example, WINNER. It is important that such tendencies occur irrespective of whether the subject can still consciously remember the earlier presentation.

With category generation the subject is asked to mention a flower or a fruit, or an object from some other category. Earlier presentation increases the tendency to mention specific items. This effect can even be observed with patients under full anaesthesia (e.g. Roorda-Hrdlicková et al., 1990). This can be demonstrated as follows. A tape is played repeating the name of one particular flower, for example, LILY or ROSE. After regaining consciousness, the patients are asked to mention any flower that comes to mind. If the anaesthesia is not too deep, a reliable tendency can be observed to mention the flower mentioned on the tape, even though patients have no conscious recollection of the episode. This finding and others indicate that attention and level of processing have little effect on the strength of encoding of implicit memories. The shifts in response tendencies also occur for unattended aspects of a message, as long as some minimal level of processing has occurred.

Effect of multiple learning trials

When a message is repeated, it is not always encoded with the same strength. A standard result in memory psychology is the lower efficiency of massed versus spaced learning trials: if messages are repeated in short succession, they are less well remembered than when repeated at longer intervals (see also Chapter 13 in this book by Neijens & Smit). An interesting study in this respect comes from the field of advertising.

Zielske (1959) presented housewives with printed advertisements according to two schedules: one group received adverts on a weekly basis (‘massed’), while the second group received adverts every four weeks (‘spaced’). Both groups were exposed 13 times. In order to obtain data about the course of recall of the advertisement over time, both groups were subdivided into smaller groups that were interviewed in specific weeks. The interviews were organized in such a way that data about both learning and forgetting between exposures were obtained (see Figure 3).

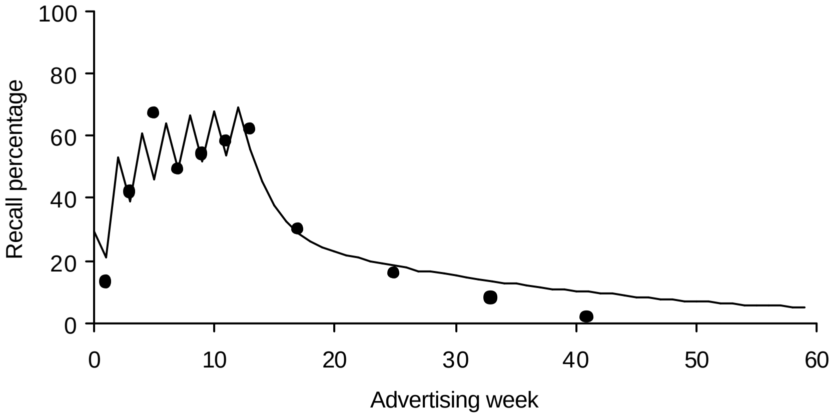
Model fit for weekly advertising schedule

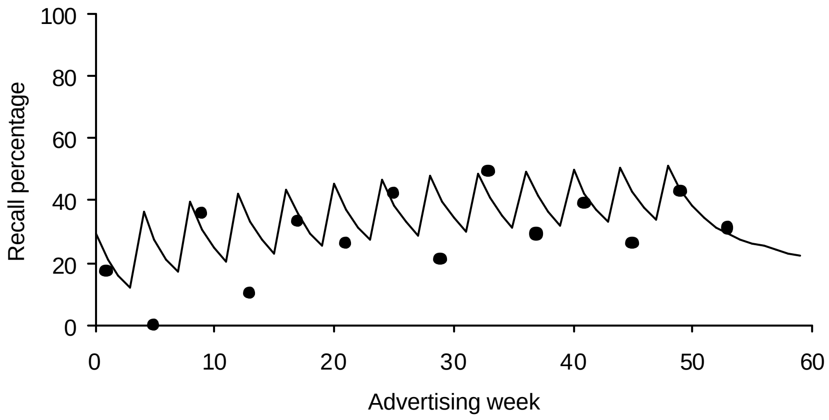
Model fit for 4-weekly advertising schedule

Figure 3: Zielske’s forgetting curve, after data from Simon (1979). The dots indicate the original data, the solid line fits by our two-store model.

The conclusions from this study were that the weekly schedule reaches a higher peak in recall, but after about 17 weeks recall improves for the spaced schedule and stays at a higher level after the completion of the campaign. Consequently, the number of 'recall weeks', that is, the number of weeks multiplied by the weekly recall rates, is larger for the four-week schedule, which implies that the decision-maker would prefer the spaced schedule to the massed one (see also Simon, 1979).

Of course, one cannot extend these conclusions beyond the domain of the data and the experimental conditions. They hold for the specific (food) product, the advertising medium, and the two exposure schedules considered. A pure data-analytic approach is unlikely to fully satisfy a decision-maker's needs. There may be other schedules that give a better performance, but this cannot be established from these data. There is some evidence that schedules which gradually lengthen the interpresentation interval, known as *expanding rehearsal* schedules, are optimal (Laudauer & Bjork, 1978). We will return to the question of optimal training schedules below, when we discuss models of memory (see also Chapter 25 in this book by Foxall).

Storage

Storage in the brain is not a passive affair. After encoding, memory processes in the brain remain active and this affects the fate of stored memories.

Short-term and long-term memory

One can view human memory as a chain of memory stores. An encoded message first passes through a short-term store that can hold a memory for a few seconds or minutes at most. It is then transferred (copied) to longer-term stores, which can last days or longer. This division in short-term versus longer-term memory was first explored systematically by Atkinson and Shiffrin (1968). They portrayed short-term memory as a 'buffer' in which only a few items can be stored. Adding a new item means dropping another from the buffer. As long as items are in the 'buffer', a process is hypothesized to be active that transfers items to long-term memory which is more resistant to forgetting. The buffer concept fits with the concept that we can only think about a few things at the same time. Baddeley and Hitch (1974) proposed a more detailed model that aimed to answer the question what working memory is used for. Their answer was: as a slave system to aid in our reasoning. We can temporarily store partial results of our thinking in such a short-term memory system and then combine them with later results to reach a final conclusion. A good example of this is mental arithmetic. They had also noticed that the buffer's capacity was less for longer words, and that it was sensitive to phonological effects. A list of syllables, BA, BI, BU, BI, was easier to remember than TA, DA, FA, DA. Baddeley and Hitch named this system the *phonological loop* because it was able to hold about two seconds' worth of speech, rather than some specified number of arbitrary 'items'. Another independent slave system was identified as the *visuo-spatial sketchpad*. Together with the central executive part of the brain, where reasoning and planning takes place, this memory model is known as the *working memory* model.

Interference and schematisation

New learning tends to interfere with existing, similar representations. Many experiments on verbal material were carried out in the early 1940s. Subjects typically studied two lists of paired associates of stimulus-response pairs such as CAR-27 and TABLE-38. One finding was that when the stimuli in the two lists are different there is little interference, but when they are similar, learning the second list can cause very rapid forgetting of the first list when the paired responses differ across lists, for example CAR-27 on the first list and CAR-84 on the second list (Osgood, 1949). In everyday life, this effect may occur when visiting two stores to compare prices over a range of products. Studying prices in the second (or third) store will interfere with your memory of prices in the first store.

When many variants of a concept (visual image or some abstract concept) are studied but the original is not, the latter is often falsely recognized as having been seen before. This has, for example, been studied with random dot patterns. The false recognition of an unstudied prototype is called *the prototypicality effect*. Connectionist models, discussed below, usually exhibit the prototypicality effect. Individually stored, similar representations may under the influence of noise spontaneously form a prototypical 'ghost representation'. Because these models mimic neural structures, it is likely that the brain also tends to form prototypes of our experiences. There are, however, other ways in which this effect may emerge, even if no prototypical 'ghost memories' are created.

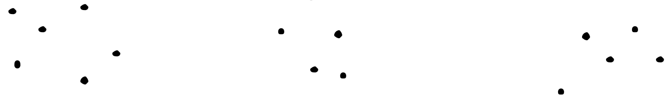

Figure 4: A prototype random dot pattern with three distorted variants. Exposure to the variants often leads to erroneous 'recognition' of the prototype.

Even without additional learning the mind appears to alter messages to fit our existing knowledge of the world. When confronted with unfamiliar material we tend to search for meaning and, when recalling the material, it will often be distorted to conform to our existing knowledge structures. Sir Frederic Bartlett was one of the first people to investigate this, in 1932. When telling a Native American traditional story to Cambridge subjects, they tended to remember a detail such as 'something black came out of his mouth' as 'he frothed at the mouth' or 'his soul expired through his mouth'. Bartlett proposed that we encode new materials in terms of *schemata* or existing knowledge structures. A consequence for the processing and remembering of advertisements, or indeed of any type of message, is that over time there will be a strong tendency towards schematisation and stereotyping. As a memory fades, the chance that the representation of the message will be adapted to fit existing knowledge increases, rather than that existing knowledge will be enhanced with the information contained in the message.

Retrieval

Accessing a memory can be as automatic as recalling our own name or as laborious as reconstructing in our mind the plot of a novel read many years ago. Retrieval is typically initiated by a stimulus that acts as a retrieval *cue*. A face may act as a cue for a name and aspects of a person. The name of a city often calls forth memories of visits. A traffic situation may cue applicable traffic rules, once learned and now remembered.

Cue effectiveness

Briefly, one could say that the better the cue, the higher the chances of memory retrieval. A study by Wagenaar (1986) illustrates this principle. For two years, he made a daily note of the most prominent event of that day, writing down its *who*, *what*, *when*, and *where*. Up to six years later he cued himself with one, two or three cues. Though the study explores many other variables, what is of interest here is that the effect of the number of cues is almost linear on the intensity of the to-be-retrieved memory (Chessa & Murre, in preparation). Intensity refers to an underlying measure of memory strength from which retrieval probability can be derived through a simple transformation (see the discussion of memory models below). Two cues are thus roughly twice as effective as a single cue.

Rather than offering more cues, taking *longer* to search one’s memory may have a similar effect when retrieving something from long-term memory. In a study by Williams and Holland (1981), adults were asked to retrieve names of schoolmates over the course of ten sessions which each lasted about an hour. The cumulative total of retrieved names increased from about 90 names in the first session to an average of 214 in the final session. It is interesting that when these data are analysed with our model (see below), retrieval time has an almost linear effect on the underlying memory intensity. This suggests that the retrieval process searches the memory space in a similar way to the way we go about finding a lost set of keys: different ‘areas’ of our memory are searched in order. Searching for two minutes is about twice as effective as searching for one minute.

Cue-induced distortions

Cues may distort to-be-retrieved memories. They may lead the interpretation of aspects of the memory. Compare, for example, How fast was the car driving when it *smashed* into the fence? with How fast was the car driving when it *touched* the fence?. Subjects will report speeds up to 15 km/hour higher for ‘smashed’ than for ‘touched’. As another example, Loftus (1977) showed how asking a question about a blue car involved in an observed accident (shown on slides) can increase errors in subsequent retrieval of the car colour. Erroneous reports of blue or blue-green cars were higher in the group who were asked the distorting question, although all subjects had seen the same *green* car.

Models of memory

Connectionist models

Connectionism (also called *neural networks*) is a modelling formalism based on the metaphor of networks of interconnected nerve cells that exchange simple signals over connections. Of particular importance is the learning capacity of many connectionist models, which makes them particularly suitable for modelling human memory. Learning is achieved by adjusting the efficiency of each connection in such a way that the behaviour of the network is slowly moulded into a type of desired or target behaviour. The target behaviour may be provided by the modeller in the form of teaching or target signals. Some types of neural networks are able to extract regularities from the stimuli to which they are exposed without being told what to aim for. They achieve this regularity learning by creating and updating internal category structures. These categories correspond to knowledge that the network has extracted from the stimuli to which it has been exposed.
Most types of neural networks are able to learn individual patterns, for example, images of faces, presented in some suitable format. When presented with an incomplete, distorted or noisy version of a pattern, networks are able to retrieve the original. Humans and animals routinely perform such pattern retrieval during cognitive processing. As outlined above, learning many versions of an unseen prototype usually leads to erroneous retrieval of the prototype in a neural network, illustrating one of many details of the ways in which neural networks approach aspects of human memory.

Recently, a number of researchers have developed connectionist models of human long-term memory, also aiming to explain certain forms of memory loss resulting from brain damage (Alvarez & Squire, 1994; McClelland et al., 1995; Murre, 1996). These models assume that two principal brain areas, the neocortex and the hippocampus, play different roles in long-term memory storage. The neocortex is the deeply folded structure covering the surface of the brain. It is viewed as the high-capacity, final repository of our memories. The hippocampus is a tube-like structure that is evolutionarily older than the neocortex. Its size is about 3% of the neocortex, making it a low-capacity system. It is, however, a much faster learner. The models cited assume that memories are first stored in—or via—the fast-learning hippocampus and are then transferred to the neocortex. The transfer of dependence from hippocampus to neocortex is called *consolidation*. While initially the retrieval of a recently experienced event is reliant upon the hippocampal system, repeated reinstatement of the hippocampal-neocortical ensemble over time results in the formation of a more permanent—hippocampally-independent—memory representation in the neocortex. The connectionist models demonstrate that such a consolidation mechanism is feasible and the simulated data resembles that of humans and animals.

A general model of learning and remembering

Chessa and Murre and Murre and Chessa (both in preparation) describe a model for forgetting and learning. The model views the memory process as a feedforward chain of memory stores. When an item is presented, a memory of the item is generated in the first store of the chain. The model assumes that during the learning process copies are generated of the memory or of a particular critical feature of the item (encoding). This generation is probabilistic.

Each copy has a probability for the time during which it is available in a store (storage). Before this lifetime expires, a copy may generate other copies in the next store in the chain (consolidation). This process of transferring representations to higher-order stores in the chain describes the increased resistance of a memory to forgetting, because lifetimes are longer on average in higher stores.

Though the memory stores in our model are kept abstract, they can be associated with neurobiological structures. In Murre et al. (in preparation), a two-store model fits data from amnesia and other memory disorders, where store 1 can be identified with the hippocampal area in the brain and store 2 with the neocortex.

In many cases, we assume that locating a single copy suffices for complete recall. We, furthermore, assume that a recall cue will typically not manage to search an entire store, but only a small, random section. This makes memory retrieval a stochastic process: even if copies are present, it is possible that none will be found, if the cued sections happen to be empty.

Learning increases the number of copies of a memory that one expects to find on average. Longer learning periods and repeated learning trials lead to a proportional increase, up to a point. When this maximum is approached, the learning process saturates and becomes less effective. Massed learning causes such saturation and is therefore less effective according to our model.

Application to learning and forgetting of advertisements

Exposure to an advertisement or similar communication generates memory copies (learning). After the exposure, their number starts to decline (forgetting). Repeated exposures simply generate more copies. In Chessa and Murre (2001) we fit our model to the Zielske data, discussed above, and to the recall of TV advertisements. Some of these results will be summarized below.

The model applied to the Zielske data contains two stores and four parameters. An exponential decay was assumed for both stores, while a saturation assumption was imposed on the learning process. The solid lines in Figure 3 show a fit of the two-store memory model. The model fits the data simultaneously, that is, the values of the parameters are the same for both schedules.

The model parameters tell us that the expected impact of the first exposure is 29%, 40% of the contents of the first store are lost every week and 4% is lost in the second, slower decaying store, while 10% of the information in the first store is rehearsed. The expected number of recall weeks for the two schedules can now be calculated exactly, which confirms the superiority of the spaced schedule by a factor of approximately 1.5.

Memory models with one and two stores were also fitted to impact data of TV advertisements, which were collected in 1997-1998 by SPOT, the Foundation for

Promotion and Optimization of Television Advertising in the Netherlands. A total of 43 campaigns were tracked for about six months. The data were obtained from 50 interviews per brand per week. The impact data are a function of TV share, measured in Gross Rating Points (GRPs) that are scheduled in the campaign weeks.

We used almost the same models as in the fits to the Zielske data, except that we did not incorporate a saturation assumption in the encoding process. The expected number of memory representations formed in the first store during exposure to an advertisement was simply taken to be a linear function of the GRPs. This means that there are three parameters for the single-store model and five for the two-store model.

The single-store model was not rejected in most cases; a two-store model was needed to fit the data of only five of the campaigns. Figure 5 shows a fit of the single-store model to impact data for a well-known brand. Its initial impact is among the highest in the tracking study. However, the forgetting rate is about 12% per week, which implies that the product needs long-term advertising in order to maintain a high impact.

The memory models not only give the advertiser a way of predicting future impact for any time-scale, but also a means to calculate the expected impact for different GRP values per week. In practice, one has the option of fixing a time-scale and a GRP budget to spend over this time-scale and then find the distribution of GRPs that yields the highest average impact. In this way, one can maximize the effect of the message to be communicated.

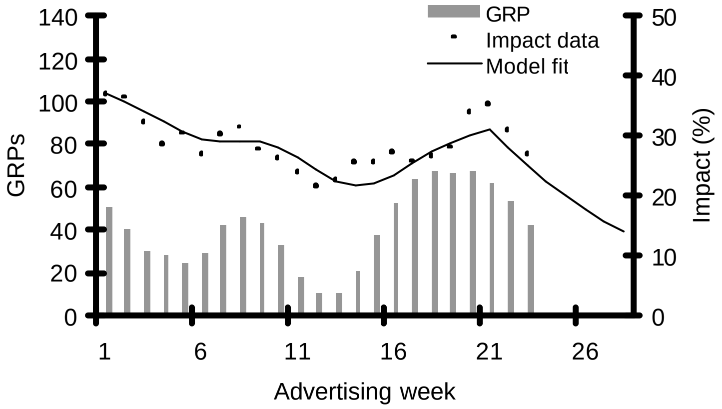

Figure 5: Impact (proven recall) of a well-known Dutch brand’s advertising strength (in gross rating points or GRPs) measured by SPOT over a period of 23 weeks. GRPs are indicated by grey bars and impact is shown by dots. The solid line is a fit of our model (single-store variant). Also shown is the predicted decline of impact seven weeks after the end of the advertising campaign.

Conclusions

It is not sufficient for a message simply to reach the intended audience. Hearing or seeing it is only the beginning of a long chain of events, with each stage feeding into the next. At each stage, essential elements of the message may be lost. Memory psychology can help to chart this chain of events. As we have discussed above, mathematical models of memory may be used to predict when the essential elements of a message will be lost, and how this can be counteracted most efficiently. The past 15 years have shown a dramatic increase in our knowledge of the biological basis of learning and memory. With connectionist models, it is already possible to investigate the complex interplay of brain and behaviour. We can expect a considerable impact of this interdiscipline of 'cognitive neuroscience' on our understanding. The influence of emotions and motivation on memory will become clearer as the underlying biological mechanisms are uncovered and the different memory 'stores' are mapped on neural mechanisms and structures. Even without this additional knowledge, however, more than a century of memory psychology has yielded a rich set of principles that can be readily applied to lead to an improvement in the efficiency of the communication of messages.

Notes

Prof. dr. Jaap M.J. Murre and dr. ir. Antonio Chessa, University of Amsterdam and University of Maastricht. Correspondence address: Jaap Murre, Department of Psychology, Roetersstraat 15, 1018 WB Amsterdam, The Netherlands. E-mail addresses: jaap@murre.com and chessa@psy.uva.nl.

References

Alvarez, R. & Squire, L.R. (1994). Memory consolidation and the medial temporal lobe: a simple network model. *Proceedings of National Academy of Sciences (USA)*, 91, 7041-7045.

Atkinson, R.C. & Shiffrin, R.M. (1968). Human memory: a proposed system and its control processes. In K.W. Spence (Ed.), *The psychology of learning and motivation: advances in research and theory Vol. 2*. New York: Academic Press, 89-195.

Baddeley, A.D. & Hitch, G. (1974). Working memory. In G.A. Bower (Ed.), *Attention and Performance VI.* Hillsdale, New Jersey: Lawrence Erlbaum, 647-667.

Bartlett, F.C. (1932). *Remembering*. Cambridge: Cambridge University Press.

Chessa, A.G. & Murre, J.M.J. (2001). A new memory model for ad impact and scheduling. *Admap, 36*(3), 37-40.

Chessa, A.G. & Murre, J.M.J. (in preparation). *A theory of learning and forgetting I: forgetting*.

Craik, F.I.M. & Lockhart, R.S. (1972). Levels of processing: a framework for memory research. *Journal of Verbal Learning and Verbal Behavior, 11*, 671-684.

Ebbinghaus, H. (1885). *Über das Gedächtnis*. Leipzig: Dunker.

Landauer, T.K. & Bjork, R.A. (1978). Optimum rehearsal patterns and name learning. In M.M. Gruneberg, P.E. Morris & R.N. Sykes (Eds.), *Practical Aspects of Memory*. London: Academic Press, 625-632.

Loftus, E.F. (1977). Shifting human color memory. *Memory and Cognition*, *5*, 696-699.

McClelland, J.L., McNaughton, B.L. & O'Reilly, R.C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*, 419-457.

Murre, J.M.J. (1996). TraceLink: A model of amnesia and consolidation of memory. *Hippocampus*, *6*, 675-684.

Murre, J.M.J. & Chessa, A.G. (in preparation). *A theory of learning and forgetting II: learning*.

Murre, J.M.J., Chessa, A.G. & Meeter, M. (in preparation). *A theory of learning and forgetting III: disorders of long-term memory*.

Osgood, C.E. (1949). The similarity paradox in human learning: a resolution. *Psychological Review*, *56*, 132-143.

Roorda-Hrdlicková, V., Wolters, G., Bonke, B. & Phaf, R.H. (1990). Unconscious perception during general anaesthesia demonstrated by an implicit memory task. In B. Bonke, W. Fitch & K. Millar (Eds.), *Memory and awareness in anaesthesia*. Amsterdam: Swets and Zeitlinger, 150-155.

Schacter, D.L. (1987). Implicit memory: history and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 501-518.

Simon, J.L. (1979). What do Zielske's real data really show about pulsing? *Journal of Marketing Research*, *16*, 415-420.

Wagenaar, W.A. (1986). My memory: a study of autobiographical memory over six years. *Cognitive Psychology*, *18*, 225-252.

Williams, M.D., & Hollan, J.D. (1981). The process of retrieval from very long-term memory. *Cognitive Science*, *5*, 85-119.

Zielske, A. (1959). The remembering and forgetting of advertising. *Journal of Marketing*, *23*, 239-243.