7.

The Architecture of Flow

Luc Steels

Nowadays many people argue that learning must be fun, but what does that mean and how can it be achieved? The fun involved in learning can surely not be like the kind of short-term 'kick' one gets from going down a roller coaster. Although this gives instant pleasure (at least to some people), true learning requires a more sustained, long-term motivational force that stimulates the learner to keep seeking new challenges and gets her through periods where the going is tough.

There are people who have a strong inner motivation to engage in challenging activities and consequently they are the ones that tend to excel. These people spend long hours in a highly focused state and seem to enjoy tremendously what they are doing. They are highly effective learners and seek out learning opportunities themselves. Typically these learners are very creative and make major breakthrough-contributions to their fields. But such focused activity is also naturally observed in children. When they are truly engaged, they strongly resist having to interrupt their play, and they generally don't like to be put in a rigid framework that constrains their activities.

We need to understand better the nature of such behaviour and how it can come about, so that we can build learning environments that stimulate self-motivated learning and teach students what the joy of learning is like. This can be achieved by a theoretical investigation in the concepts underlying self-motivated learning and development, and by looking at case studies where high states of motivation are clearly reached by students and teachers alike.

Such case studies can be found in some pre-school environments, like in Reggio Emilia (Rinaldi,2003), in successful art academies and design institutes, like the Antwerp Fashion Academy or the Ivrea Design Institute, in innovative business schools, and even in some enlightened science universities where students are immersed in laboratory practice and work intensively in group towards the realisation of particular projects. Unfortunately, most schools and curricula in contemporary societies, particularly at the level of higher education, do not exhibit the necessary characteristics to generate sustained self-motivated learning.

This paper first surveys discussions on motivation in psychology and its impact on educational practice. Then I report briefly on attempts to make cognitive science models of ‘autotelic agents’ which have an inner drive towards self-development. The paper concludes with some reflections on what the implications for education could be.

2. Theories of Motivation

Reinforcement Learning

Most models in psychology and neuroscience today are still rooted in the behaviourist framework of reward and punishment, originally coming from the work of Skinner and his associates (Skinner,1953). Consequently the pedagogical strategies that we find in contemporary schools are still to a large extent based on these ideas. There are three major assumptions:

First, behaviourism argues that the environment (or a trainer) can trigger adaptation and learning by providing rewards, for example food or other means that give direct pleasure, or punishments, for example through the inducement of corporal pain. Rewards reinforce specific behaviours because they inform the individual that they are beneficial, in other words that a viable state can be reached and maintained. Punishments signal that the behaviours that were enacted need to be abandoned or new knowledge and skills need to be acquired.

Second, classical behaviourism proposes that learners start with a repertoire of reflex behaviours and an innate value system. New behaviours are shaped by reward and punishment produced by the environment. When a trainer or educator hands out the reward or punishment, she can push development in specific directions and the trainer's value system may become progressively internalised by the trainee. This opens up the possibility of cultural transmission of value systems.

Third, classical behaviourism proposes that this reinforcement framework is an adequate theory of action selection and motivation, in the sense that the main purpose of an individual is to seek reward and avoid punishment, and all the rest (acquisition of new behaviours and internalisation of a value system) follows.

Behaviourism has had a tremendous impact both on education and on behavioural and brain sciences. In the educational domain, it has lead to an emphasis on disciplinary relations between teachers and students where fear for punishments and exams are the primary instruments for motivating students. Technologically, it has lead to so called ‘programmed instruction’ systems, as practiced for example in ‘language laboratories’. In the behavioural and brain sciences, it has given rise to a tradition where results from experiments on rats are routinely transferred to humans and where representation making, meaning, personality development, etc. are taboo subjects.

Criticism of the reinforcement paradigm

In the past decades, there has been a growing feeling that the behaviourist approach is not a very 'humane' way to organise education, even though it can sometimes give very good results for the acquisition of specific skills. Moreover several educational thinkers, such as Illich(1973) or Postman(1972), have pointed out from the early seventies onwards that there are severe disadvantages using this approach in the long term, particularly for educating children to become responsible, autonomously acting members of the community. They argued that this approach to education stifles creativity and personality development.

They proposed an alternative approach: to give pupils more freedom, abandon systematic grading, and frame the teacher in the role of supporting the learning process rather than being the source of reward and punishment. The pupil is encouraged to explore domains of knowledge by herself, instead of being supplied with pre-established answers. Specialised schools, like the Freinet schools (Sivell,2004), the Reggio experiment in Italy (Rinaldi,2003), or the Sri Atmananda Memorial schools in India are examples where this approach has worked - albeit usually only for pre-school or the first grades in primary school.

But this alternative learning model has not scaled very well to the whole school system and is now regarded as inapplicable on a large scale. Where it has been tried beyond primary school, academic achievement of students are plummeting and teachers are confronted with pupils that lack the motivation to learn, resulting in boredom and subsequent violent or deviant behaviour. It seems that the utopian models proposed in the nineteen seventies only work when the student body is very homogenous, the community, including the parents, fully co-operate, and the non-school environment (particularly the leisure activities) do not overpower children as they do in our contemporary society.

In a recent essay that was distributed to all teachers in France, the French minister for education Luc Ferry (2003) blamed the current educational crisis in France to a breakdown of the traditional model, with the teacher as authority and transmitter of knowledge. He attributed this breakdown to the propagation of ideas related to 'free' alternative schooling and suggested a return to the traditional teacher-centred model. French teachers welcomed his proposals with public burnings of the book and month-long strikes! The latter was actually a symptom of a deeply rooted crisis in the teaching profession, partly caused by an abandonment of the traditional role of the teacher. In other words, the teachers are as much lost as the students.

3. Flow theory

So how can we go beyond the traditional reinforcement learning framework without falling into the other extreme of a completely unstructured free environment? The theory of self-motivated learning and behaviour may give us a potential clue. This theory was originally developed by the humanistic psychologist Csikszenmihalyi, based on studying the activities of painters, rock climbers, surgeons, and other people who showed to be deeply involved in some very complex activity, often for the sake of doing it, i.e. without direct reward in the form of financial or status compensation (Csikszentmihalyi,1978). He called these activities autotelic. "Autotelic" signifies that the motivational driving force ("telos") comes from the individual herself ("auto") instead of from an external source, administered by rewards and punishments.

Autotelic activities induce a strong form of enjoyment, which Csikszentmihalyi has characterised as "flow". The word "flow" is a common sense word and so there is a risk to interpret it too broadly. Csikszenmihalyi intends a restricted usage, being a state which often occurs as a side effect of autotelic activities:

> People concentrate their attention on a limited stimulus field, forget personal problems, lose their sense of time and of themselves, feel competent and in control, and have a sense of harmony and union with their surroundings. (...) a person enjoys what he or she is doing and ceases to worry about whether the activity will be productive and whether it will be rewarded. o.c. p. 182.

Because the activity is enjoyable, the person who experiences this enjoyment seeks it again, i.e. it becomes self-motivated. Moreover due to the high concentration and the strong self-motivation, learning takes place very fast. The learner is eager to find the necessary sources and tools herself and spends time on the acquisition of skills, even if they are not exciting in themselves, as long as they contribute to the autotelic activity.

Given this definition of flow, it is quite obvious that many people will have experienced some form of flow in their life, and that children in particular enter into flow experiences quite often, particularly during play. Flow is sometimes associated with the ultimate high experience of the rock climber that has finally managed to climb Mount Everest, but that is an exceptional situation. Flow - as defined here - is much more common and can just as well happen in every-day experiences like playing with children or engaging in a long-term love relationship.

It is also important to distinguish flow from directly pleasurable activities like going down a roller coaster. A key difference is that the activity must in itself be challenging - otherwise there is no feeling of satisfaction after difficulties have been surmounted. Moreover there must be a steady progression in the nature and particularly the level of the challenge. This is the reason why child rearing can be so enjoyable and fascinating. A child keeps developing all the time - which is what makes the interaction fun - and that creates continuously new challenges for the parent to figure out what she is thinking, what she might want to do or not do, and so on. The rock climber can also scale up the level of difficulty with which rocks are being climbed or the kinds of rocks that are tackled. Similarly, the musician can first play easy pieces and then steadily move up. She can first play with other amateur musicians and then play with better and better musicians. The performance can be first for a few friends, but then for a larger and larger unknown audience.

An obvious key question is: What makes activities autotelic? This is where Csikszenmihalyi's most important contribution comes in. He argues that it lies in a good balance between high challenge, generated through the activity and perceived as meaningful to the individual, and high skill required to cope with this challenge:

> Common to all these forms of autotelic involvement is a matching of personal skills against a range of physical or symbolic opportunities for action that represent meaningful challenges to the individual. o.c. p. 181

When the challenge is too high for the available skill or the opportunity for action so bewildering that no clear course can be seen, anxiety sets in, particularly when there is no hope to develop appropriate skills by learning. The person gets paralysed and eventually may develop symptoms of withdrawal and depression. When the challenge is too low for the available skill, boredom sets in and the long-term reaction may be equally negative. So the optimal regime is somewhere between the two, when there is a match of challenge and skill (figure 1).

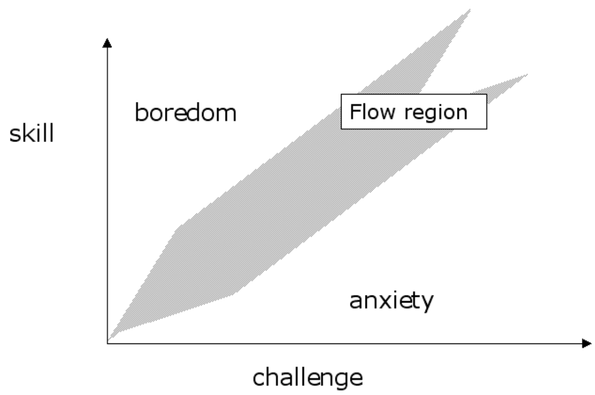

Figure 1. The flow region is defined by a balance between challenge and skill. When challenge is too low for the skill level reached, there is a feeling of boredom. When skill is too low with respect to the level of challenge, there is a feeling of anxiety.

The individual needs the ability to decrease challenge when this become too high, or alternatively be given the time to build up more skill to cope with the challenge. But it is equally important that the individual is in control of the challenge level, so as to increase challenge when the skill has become higher, otherwise activities become boring and self-motivation is lost.

Flow is not a steady state but occurs in the process of reaching a balanced state with high challenge and skill. This balanced state never lasts. When an activity is carried out often, skill normally increases so that the activity becomes boring and new challenges need to be found. So the individual seeking the flow experience is always 'on the move'.

Implications of the Flow concept

The insights of Csikszenmihalyi and those who have elaborated these insights further (see for example papers referenced in Csikszentmihalyi and Selega(2001)) are potentially of enormous significance. They hint to an alternative to classical behaviourism and hence to a very different approach to psychology and neuroscience compared to the one suggested by the behaviourist paradigm. They also complement the information processing approach to psychology, which has tended to ignore issues related to emotion, motivation and affect. The implications for education are very deep.

The current school system, in as far as it is built on the concept of reward and punishment, does not create the conditions for flow, in fact it often totally ignores them. The meaningfulness of an activity is not considered as primary (often it is not even meaningful to the teacher). Instead activities are imposed in a top-down fashion and based on statewide uniform programs, such as the so-called 'collège unique' program in France. The learners have no individual control over the challenge levels in an activity. This results in giving up any effort to learn as soon as they see no way to cope, or to get bored and become totally disinterested in the school as a whole when the challenge is not high enough.

A typical example of a traditional teaching situation where the conditions for flow are destroyed, has been described by John Matthews, who is an authority in visual art education for children (Matthews,1993). In a provocative essay entitled "Art Education as a form of child abuse", he argues that the natural desire and enjoyment that children experience in developing their visual skills is totally ignored by the educational system. Specifically in the case of Singapore it even becomes completely wiped out so that children lose all interest in the matter:

Generally, from the primary level onward, rigid, highly prescriptive approaches to drawing are employed. Typically, the teacher instructs, in a series of strict stages, toward a pre-envisaged, fixed end-point, usually of the most stereotypical, banal kind. The children are required to obediently follow every step, and if variation is allowed at all, it is of the most trivial kind. Identical art works are the result. If a child has varied a little from the fixed procedure, he or she is often criticised, if not punished. Many recorded observations testify that, out of a class of forty children, the one or two art works which possess any vitality are condemned and their makers accused of an inability to concentrate or follow instructions. Other recordings made in private art classes for the very young show adults sometimes altering their children's drawing or even re-drawing the image wholesale. The drawing process is often totally dominated by the adult, who gives strict instructions about the sequence of strokes, from which any variation is strictly criticised. Teachers sometimes even physically manipulate the child's drawing arm itself.

By the time children get to the end of primary school, they have lost all interest in drawing or painting and are visually illiterate. The culprit is obviously that goals are imposed in a top-down fashion through reward and punishment instead of being generated in a self-motivated way.

The 'free school' paradigm proposed by Illich and others goes a step in the right direction because it accepts that the learning environment must be adapted to the individual student and that the learner must have control to some extent about what should be learned and when. But at the same time, this paradigm equally ignores the necessity for the teacher and the school environment to create scaffolded opportunities for action that generate the challenges for self-motivated learning.

So two things are needed. A learner must be able to feel some control of the challenge level, but at the same time, the environment is crucial in generating new opportunities and providing structure to the learning experience. For example, the rock climber must be able to choose a more difficult or more challenging rock but there must of course be rocks in the first place. The rocks are the learning environment for the climber, a continuing source of increasingly more challenging opportunities for enhancing the skill. In the same way, the school, the teachers, and the parents must create the learning environment, which must be full of opportunities, challenges, but also the resources to progressively cope and grow in the acquisition of a particular skill.

According to this viewpoint, the most important thing that children have to learn in school is not so much knowledge or skill about a particular subject matter (for example, copying the drawing of the teacher in exactly the same way), but rather the experience that there are such things as autotelic activities, which are not only highly rewarding on a long-term time scale at a deep level, but also a foundation for self-actualisation.

The flow concept is also relevant for the teachers: Teaching can be a highly rewarding activity, when there is a smooth dynamical interaction between teacher and students and students progressively become better in a certain domain of expertise under the guidance of the teacher. But this requires that the teacher must have a way to control the situation enough to generate the flow experience. However, a teacher won't get a grip on the teaching process and motivation will break down if a class contains too many students of different levels and backgrounds, if the students are not at all interested to engage themselves in the learning process, or if the content changes too fast and is imposed top-down by administrations.

The school must be equally rewarding for the teacher as for the student. If the teachers are not enjoying what they are doing, then they are not going to be able to sustain appropriate learning environments either. It has been argued that interference from political and bureaucratic institutions that impinge on the freedom of teachers to set the course of action is one of the reasons why teachers become anxious or bored with their jobs. The student-teacher dynamic has been broken in contemporary schools from both sides (Nevejan,2003).

Impact

The first obvious way in which the ideas associated with flow theory can impact education is to propagate them widely in the teacher community and among politicians or administrators that set or enact educational policy. This work has been taken up by Csikszentmihalyi in a series of popular publications, such as (Csikszentmihalyi,1990).

The second approach is to engage in a scientific study of flow. This challenge has also been taken up by several researchers, including Antonella Delle Fave(2004). Most of this research at the moment is observational in nature; it tries to find out how people characterise self-motivated activities, when they have them, which groups have them more than others, how they are triggered, etc. This work is of course extremely useful to get an empirical grip on the phenomenon, but there is also another approach possible, which is to develop operational models of the cognitive mechanisms that underlie autotelic behaviour.

Traditional symbolic cognitive models developed in AI (Newell,1990), as well as behaviourist and neo-behaviourist models, such as connectionism (Churchland and Sejnowski,1995) have well worked out formalised theories of the architecture of a 'cognitive agent'. This has made it possible to define and study systems based on reinforcement learning (Sutton and Barto,1998) and even to use these models in artificial devices like robots that are acquiring crucial features of the world and how to behave in it (Steels and Brooks,1995), (Pfeifer, 2000). Recent work (Steels, 2004) has shown that it is in fact possible to develop mathematical and computational models for autotelic behaviour and instantiate them in self-developing robots. This research can tell us something about the kinds of internal mechanisms that are required to induce autotelic behaviour and how they work. It can then be used to stimulate the development of these mechanisms in human children or to incorporate them in learning environments.

Yet another approach (explored by François Pachet (2004)) is to build devices that elicit flow in users. One example he discusses is the `Continuator', a music system that acquires the statistical properties of musical input introduced through a keyboard and then mirrors this with its own output which consists of variations on the learned patterns. Experiments with human subjects show that this simple principle (when implemented well) generates enormous excitement in children and adults alike. Because the system continuously learns, the interaction remains always novel and users try to elicit more and more complex behavior. Building this kind of `flow machine' (in the sense of machines that help to generate flow) has a dual purpose: They can tell us something about why humans become very excited in interacting with such devices, and also what features a system needs to have in order to elicit flow.

In my own work, I have focused on the second approach, i.e. trying to make operational models of agents that simulate aspects of autotelic behaviour. The next section reports briefly on the kinds of models that I have been able to construct so far.

4. Autotelic Agents

The models of agent architecture that cognitive scientists have attempted to construct involve a large number of components: for sensory perception, motor control, categorisation, memory, situation recognition, action planning, action selection, motivation, communication, etc. Each of these components takes some input (for example a bitmap produced by a camera in the case of a vision system) and maps this onto some output (for example an analysis of the scene). Although earlier versions of agent architectures did not take learning into account, it is now common to make every component in the agent's architecture adaptive: new low level visual feature detectors arise, new actions are learned, new plans are made and stored for later, new perceptual categories are acquired, new conventions for communication invented or adopted, etc.

Today we have already a vast range of quite detailed operational models for many of the components and learning strategies for these components (Pfeifer and Scheier, 2000). Some are argued to be realistic models of the brain. The learning strategies are either unsupervised, meaning that they attempt to find recurrent structures in large data sets using statistical methods, or supervised, meaning that they rely on a re-enforcement signal coming from the environment or from a teacher. Supervised learning mechanisms operationalise the behaviourist re-enforcement framework discussed earlier in the paper.

In all proposals so far, there is however no overall motivational mechanism that is helping to decide what to do, while pushing the agent towards greater levels of achievement. Many agents disussed in the literature are purely reactive, responding to environmental stimuli – which might be a good model for an ant but not for a human person. If any kind of development is modeled, it is completely steered by human designers with carefully scaffolded input data and reward functions. Learning of a particular task is usually studied in complete isolation, for example, the researcher focuses on the acquisition of new visual features and supplies input-output pairs focusing exclusively on this task.

To make cognitive agents that are autotelic, I argue that four ingredients need to be added to the agent architectures proposed so far in the literature:

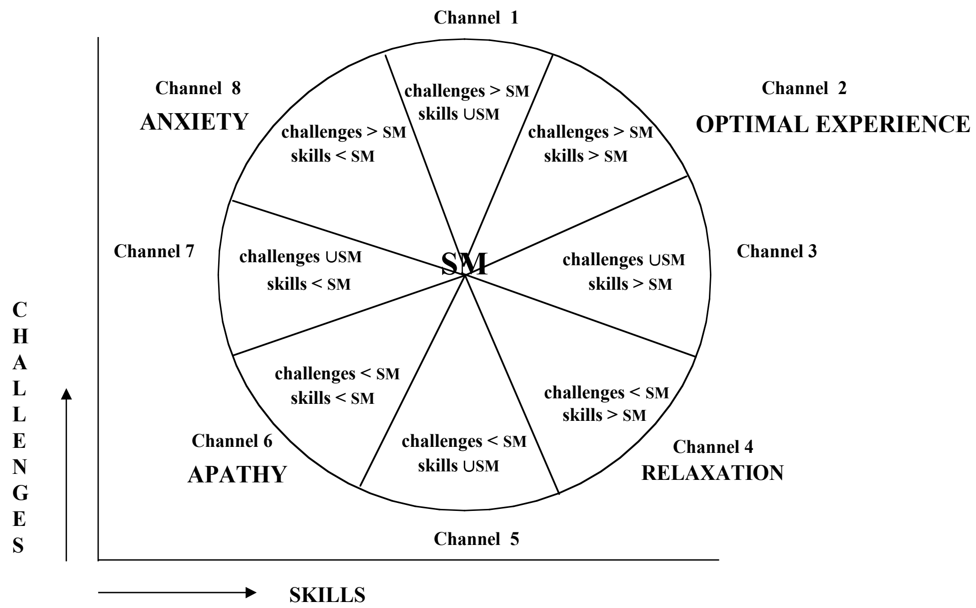

Figure 3. The quality of experience can be defined as a function of challenge and skill levels. They are compared with the standard subjective mean (SM) skill and challenge levels experienced by the agent in the past (see Delle Fave, this volume)

1. The agent needs a way to sense the quality of experience, in the sense of Csikszentmihalyi (1990) (see figure 3). Because quality of experience is a function of skill and challenge, it follows that the agent needs internal sensors that are capable to track skill levels, and categorisation processes that estimate the degree of challenge associated with a particular situation in the environment. Because quality of experience is based on a comparison to a subjective mean (SM) of skill and challenge levels, it follows that challenge and skill levels need to be tracked over longer periods of time. Clearly developing sensors and categorisations of skill and challenge is not a static thing, but must be learned and continuously adapted based on feedback of performance.

2. The agent needs a way to increase or decrease the challenge level of a task. This can be done by seeking out (or avoiding) environments or tasks that have more (or too much) complexity, or by regulating internally with what degree of complexity environmental stimuli should be treated and with what precision actions executed. For example, a vision system might perform a very detailed scene analysis, recognising all objects in the scene and their interrelationships (in other words set a very high challenge level), or alternatively it may perform a simple rough scene analysis with only vague contours of objects and no further recognition (implying a very low challenge level).

3. The skill levels of the agent change due to learning. So if a situation is encountered frequently, skills are exercised and skill levels automatically increase. Alternatively, if a skill is not exercised, it deteriorates, causing the skill level to decrease. (Compare this to experience in learning a foreign language. Knowledge of the language rapidly increases when one speaks it a lot but it decreases just as quickly when one no longer has the opportunity to speak.) Exercising a skill requires psychic resources (memory,

processing time), and this may be under partial control of the agent, so that skill levels can effectively be decreased causing the agent to pay less attention.

4. In most action selection models so far, the criteria for choosing an action are uniquely based on the utility of the action towards achieving specific goals. Action is here taken in a broad sense and also includes perception, i.e. attending to a particular object in the environment. In the case of autotelic agents, optimising the quality of experience becomes an additional criterion for action selection. In other words, an action that has less utility for the agent, but has greater value in optimising the quality of experience, might still be preferred. This principle needs to operate on different time-scales: Long term benefit is larger if a steep slope of skill increase has to be achieved to cope with a particular challenge because that will enable future experiences with higher skill and higher challenge.

5. Our experiments have shown that autotelic agent need occasionally to `shake up' the challenge levels they have set themselves, to avoid getting into locally optimal states which miss out on opportunities to reach globally higher skill and challenge levels. In other words, the agent cannot always rely on the environment to pose new opportunities and challenges, but must seek them out, thus generating an ever increasing self-developmental spiral of higher challenge followed by matching skill.

'Flow' is usually characterised as a state in which there is a high concentration and sharp attention to a specific task and aspects of the environment relevant for that task. In the agent architecture discussed here, attention is operationalised by a focusing of mental resources (time, memory) and an increase of challenge levels for the relevant components or tasks (with subsequent decrease of irrelevant components). Flow automatically develops when the agent self-regulates challenge and skill in order to maximise the chance that optimal experience might occur. So flow is an emergent side effect of the model and not a causal factor in itself.

Once a model of a single autotelic agent exists, it is possible to experiment with a population of autotelic agents which each have their own developing skill and challenge levels. Often one agent must lower his challenge levels, and hence the complexity of his behavior, because the agent with which he is interacting is not yet able to deal with this complexity. On the other hand, an agent that is interacting with another agent that has higher challenge levels, needs to increase his skills to cope with the interaction. So agents coordinate their challenge levels and can push each other up towards reaching higher challenges, and hence opportunities for reaching higher skills.

6. Implications for education

This is not the right place to go into further technical detail of these operational models, nor discuss simulation results (see Steels, 2004). Instead we turn to the question of what we can learn from these models with respect to bettering human learning environments. I believe the implications are quite straightforward and essential for the practical implementation of learner-centered learning:

1. It appears essential that students develop an awareness of their own skill levels. In some sense exams and tests already do this. But often tests are seen as hurdles that need to be overcome, and once they are overcome, everybody in the same class is seen as being on

the same level. This is clearly not valid. Every student has their own skill levels which are dynamically changing, and tests should be presented to the student as a tool to monitor their skill levels.

2. Students should develop competence in estimating the challenge level in a particular task situation. Usually the teachers scaffold challenge, for example by first giving simple exercises and then gradually providing more complex ones. This is in fact one of the major tasks of the teacher, and Vygotsky's notion of the zone of proximal development captures this idea very well. But in real life, one must be able to decide yourself when a challenge is appropriate for your skill, and this implies that one must develop appropriate categories for judging challenge levels.

3. Students should be stimulated to develop an increased awareness that skill can be increased through learning and that skill needs to be maintained by continuously exercising it. The motivation for learning has to come from the realisation that increases of skill are going to bring a whole range of new challenges within grasp. They should be able to see that a particular problem or task is too difficult right away but that work on a supporting skill can help to achieve the right level.

4. Finally the learner needs to be given control of the challenge level, perhaps not all the time because then there is a risk, particularly with less ambitious students, to remain stuck with skill levels that are below the ultimate capacity of the student. But, as reported by Csikszentmihalyi in many of his case studies, self-control of the challenge level is one of the absolute prerequisites for reaching a flow experience. Self-motivation dies out if somebody is felt to have no control of his destiny.

It is obvious that many curricula and learning environments in one way or another incorporate some of the principles expressed here, thanks to the intuitions of good teachers. But just as many school environments violate the principles suggested here. The more we learn about autotelic experiences and how they come about, the better we can tailor learning environments to exploit the natural tendencies of human beings to seek optimal experience and balance challenge and skill through self-motivated learning and self-control. In many cases, it is a matter of giving learners much greater control in setting challenge levels, deciding what resources to bring to bear, seeking out environments in which supporting skills might be acquired, or interacting with peers that have slightly higher skills.

5. Conclusions

Flow is part of a long-term engagement, in which the individual is continuously trying to go beyond existing skills to reach new heights. Only then the deep experience of satisfaction can occur that sustains new learning. When a school curriculum can be organised in such a way that students feel progressively challenged – without experiencing the anxiety of overreaching their skill – very powerful natural learning processes get triggered and students become self-motivated.

This paper briefly reported on efforts to make operational models of 'cognitive agents' that incorporate some of the intuitive ideas associated with flow. These models include components for sensing skill and challenge levels and determining the quality of experience, mechanisms to increase or decrease challenge and skill levels and thus focus physic energy, and action selection mechanisms that take optimal experience into account in addition to mere utility with respect to short-term goals.

Much remains to be discovered about the nature of autotelic activities and the design and experimentation with cognitive architectures that integrate notions of flow is a very new area, but I hope that the ideas expressed in this paper make it clear that progress is being made and that the potential impact on the future of learning is high.

Acknowledgements

I thank all the participants from the various workshops on the 'Future of Learning' for many stimulating thoughts and comments, and Mario Tokoro for continuing encouragement to pursue these topics in the best circumstances imaginable at the Sony Computer Science Laboratory in Paris. Discussions with Frédéric Kaplan and François Pachet have helped to fine-tune the agent architectures of flow which form the background to this paper. I also thank Marleen Wynants for her critical comments on the notion of flow in general and on the present paper in particular.

References

Churchland, P. and Sejnowski, T. (1995) The Computational Brain. The MIT Press, Cambridge Ma.

Csikszentmihalyi, M. (1978) Beyond Boredom and Anxiety: Experiencing Flow in Work and Play. Cambridge University Press, Cambridge.

Csikszentmihalyi, M. (1990) Flow. The Psychology of Optimal Experience. Harper and Row, New York.

Csikszentmihalyi, M. and I. Selega (Editors) (2001) Optimal Experience: Psychological Studies of Flow in Consciousness. Cambridge University Press, Cambridge.

Delle Fave, A. (2004) A Feeling of Well-Being for Learning and Teaching. This volume.

Elman, J. (1993). Learning and development in neural networks: The importance of starting small. Cognition, v. 48. p. 71-89

Ferry, L. (2003) Lettre a tous ceux qui aiment l'école. F. Jacob, Paris.

Hedegaard, M. (2003) In: Tokoro, M. and L. Steels (eds.) The Future of Learning. Issues and Prospects. IOS Press Amsterdam, 2003. p. 79-92.

Illich, I. (1973a) Deschooling Society, Harmondsworth: Penguin.

Johnson, M. (1996) The Infant Brain. In: Tokoro, M. and L. Steels (eds.) The Future of Learning. Issues and Prospects. IOS Press, Amsterdam. p. p. 101-116.

Matthews, J. (1993) Art Education as a form of child abuse. Lecture for the National Institute of Education Singapore.

Nevejan, C. (2003) Integrated Learning Environments. In: Tokoro, M. and L. Steels (eds.) The Future of Learning. Issues and Prospects. IOS Press, Amsterdam. p. 33-49.

Newell, A. (1990) Unified Theories of Cognition. Harvard University Press, Cambridge, MA.

Papadimitrious, C. and K. Steiglitz (1998) Combinatorial Optimization: Algorithms and Complexity. Dover, New York.

Pfeifer, R. and C. Scheier (2000) Understanding Intelligence. The MIT Press, Cambridge Mass.

Postman, N. (1972) The End of Education. Redefining the value of school. Vintage Books, New York.

Rinaldi, C. (2003) The Joys of Pre-school learning. In: Tokoro, M. and L. Steels (eds.) The Future of Learning. Issues and Prospects. IOS Press, Amsterdam. p. 57-78.

Sivell, J. (2004) Tools for Embodied Teaching: Célestin Freinet and the Learner-Centered Classroom. (This volume)

Skinner, B.F. (1953). Science and Human Behavior. Macmillan, New York.

Steels, L. and R. Brooks (eds.) (1995) The Artificial life Route to Artificial Intelligence. Building Situated Embodied Agents. Lawrence Erlbaum, New Haven.

Steels, L. (2003) Social Language Learning. In: Tokoro, M. and L. Steels (eds.) The Future of Learning Issues and Prospects. IOS Press, Amsterdam. p. 133-159.

Steels, L. (2004) Autotelic Agents: Robots that have fun. In: Hafner, V., F. Iida, R. Pfeifer and L. Steels (eds.) (2004) Embodied AI. Proceedings of the Dagstuhl Seminar. Lecture Notes in Computer Science. Springer Verlag, Berlin.

Sutton, R. and A. Barto (1998) Reinforcement learning: an introduction. MIT Press, Cambridge Ma.

Thelen, B. and L. Smith (1994) A dynamic systems approach to cognition and development. MIT press, Cambridge Ma.

LUC STEELS is director of the Sony Computer Science Laboratory in Paris and professor of Artificial Intelligence at the University of Brussels (VUB). He was educated at the University of Antwerp (Belgium) and MIT (USA). His research has covered a broad range of interests in autonomous cognitive robotics and artificial life. Steels currently researches meaning and the origins and learning of language through computational and robotic models. He has developed the framework of evolutionary language games as a way to study the acquisition of concepts and language grounded and situated in experience. The framework allows a systematic study of transactional learning and could lead to novel learning environments for language acquisition. Steels published and co-edited a dozen books on various topics in artificial intelligence and regularly publishes in top scientific journals, with recent articles in Behavioral and Brain Sciences, Proceedings of the Royal Society, Robotics and Autonomous systems, Artificial Life, a.o. He is a member of the New York Academy of Sciences and a fellow of ECCAI, the European AI organisation.