

Art & Science as By-Products of Search for Novel Patterns, or Data Compressible in Unknown Yet Learnable Ways

Jürgen Schmidhuber

IDSIA, Galleria 2, 6928 Manno-Lugano, Switzerland

University of Lugano & SUPSI, Switzerland

TU München, Germany

juergen@idsia.ch - <http://www.idsia.ch/~juergen>

Abstract

I argue that science, art, music, comedy are just by-products of our intrinsic desire to create / discover more *novel patterns*, that is, data compressible in hitherto unknown ways. It is possible to rigorously formalize this concept and implement it on learning machines, thus building artificial robotic scientists and artists equipped with intrinsically motivated curiosity and creativity. I summarize our work on this topic since 1990; for concrete implementation details (1990-2009) see [17, 16, 30, 20, 21, 24, 25].

This overview heavily draws from previous papers for a more science-oriented audience [17, 16, 30, 20, 21, 24, 26, 25]. In fact, certain parts of it are just minor revisions of previously published work [25, 26] based on earlier papers [21, 24, 19] and thus are ©Springer-Verlag Berlin Heidelberg 2009 and the Society of Instrument and Control Engineers (reprinted with permission).

Contents

| | | |
|----------|---|-----------|
| 1 | Novel Patterns Through Compression Progress Drive | 2 |
| 1.1 | Patterns and Subjective Simplicity or Beauty | 3 |
| 1.2 | Novel Patterns and Subjective Interestingness | 4 |
| 1.3 | Maximizing the Joy of Finding Novel Patterns | 5 |
| 1.4 | Relation to External Reward | 5 |
| 2 | Consequences of Novel Pattern Search | 6 |
| 2.1 | Art & Music as By-Products of Novel Pattern Search | 6 |
| 2.2 | No Objective “Ideal Ratio” Between Expected and Unexpected | 7 |
| 2.3 | Difference to Traditional Information Theory | 8 |
| 2.4 | How <i>Active</i> Artists and <i>Passive</i> Observers of Art are Alike | 8 |
| 2.5 | Science as By-Product of Novel Pattern Search | 8 |
| 2.6 | How Artists and Scientists are Alike | 9 |
| 2.7 | Jokes and Comedy and Other Sources of Fun | 9 |
| 3 | Concrete Implementations of Systems that Find or Create Novel Patterns | 10 |
| 3.1 | Reward for Prediction Error | 10 |
| 3.2 | Reward for Predictor Improvements | 10 |
| 3.3 | Relative Entropy between Agent’s Prior and Posterior | 11 |
| 3.4 | Learning Programs & Skills Through Zero Sum Reward Games | 11 |
| 3.5 | Optimal Systems Driven by Prediction Progress & Compression Progress | 12 |
| 3.6 | No Essential Difference to Human Artists & Scientists? | 12 |
| 3.7 | Improving Real Reward Intake | 13 |
| 3.8 | Other Implementations | 13 |
| 4 | Illustrations of Subjective Beauty and its <i>First Derivative</i> Interestingness | 13 |
| 4.1 | A Pretty Simple Face with a Short Algorithmic Description | 13 |
| 4.2 | Another Drawing That Can Be Encoded By Very Few Bits | 14 |
| 5 | Conclusion & Outlook | 14 |
| 5.1 | Acknowledgments | 15 |

1 Novel Patterns Through Compression Progress Drive

Karl Popper claimed that “all life is problem solving” [15]. To solve existential problems such as avoiding hunger or heat, a baby has to learn how the initially unknown environment responds to its actions. To find mother’s breast whenever hunger sensors get active is a non-trivial action sequence involving numerous neural muscle control signals in response to tactile and other sensory stimuli.

When there is no immediate need to satisfy thirst or other built-in primitive drives, the baby does not run idle. Instead it actively conducts experiments: what sensory feedback do I get if I move my eyes or my fingers or my tongue just like that? Being able to predict effects of actions will later make it easier to plan control sequences

leading to desirable states, such as those where heat and hunger sensors are switched off.

The growing infant quickly gets bored by things it already understands well, but also by those it does not understand at all, always searching for new effects exhibiting some yet unexplained but learnable regularity. It acquires more and more complex behaviors building on previously acquired, simpler behaviors. Eventually it might become a physicist discovering previously unknown physical laws, or an artist creating new eye-opening artworks.

For a long time I have been arguing, using various wordings, that all this behavior is driven by a very simple mechanism that maximizes the internal joy or intrinsic reward for the discovery of *novel patterns*. Both concepts are essential: *pattern*, and *novelty*. Let us nail them down formally such that we can implement this drive on computers and robots, to build artificial scientists and artists.

1.1 Patterns and Subjective Simplicity or Beauty

What exactly is a pattern? Computer science has the answer: Any set of raw data, such as the history of some observer's previous actions & sensations & rewards, exhibits a pattern or regularity if there is an algorithm or computer program that can compress it [11]. Random noise is irregular and arbitrary and incompressible. But videos are regular as most single frames are very similar to the previous one. By encoding only the deviations, movie compression algorithms can save lots of storage space. Complex-looking fractal images [12] are regular, as they usually look a lot like their details, being computable by very short programs that re-use the same code over and over again for different image parts. The entire universe is regular [23]: every photon behaves the same way; gravity is the same on Jupiter and Mars, mountains usually don't move overnight but remain where they are, etc.

There is no general way of finding the shortest program computing any given data [11]. This shall not worry us. Let $O(t)$ denote the state of some subjective observer O at time t . O has some method for compressing the growing data history. According to our *lazy brain theory* [19, 18, 21, 24, 26, 25] we identify the subjective simplicity or compressibility or regularity or beauty $B(D, O(t))$ of a new observation sequence or action sequence D (but not its interestingness or aesthetic / artistic value - see below) as the negative number of bits required to encode D , given the observer's current limited prior knowledge and limited compression method.

For example, to efficiently encode previously viewed human faces, a compressor such as an artificial neural network may find it useful to generate the internal representation of a prototype face. To encode a new face, it must only encode the deviations from the prototype [19]. Thus a new face that does not deviate much from the prototype will be subjectively more beautiful than others. (This immediately explains why many human observers prefer faces similar to their own—what they see every day in the mirror will influence their subjective prototype face, for simple reasons of coding efficiency). Similarly for faces that exhibit geometric regularities such as symmetries or simple proportions [24]—in principle, the compressor may exploit any regularity for reducing the number of bits required to store the data.

Generally speaking, among several sub-patterns classified as *comparable* by a given observer, the subjectively most beautiful (but not the most interesting or aesthetically rewarding - see below) is the one with the simplest (shortest) description, given the observer’s current particular method for encoding and memorizing it [19]. For example, mathematicians find beauty in a simple proof with a short description in the formal language they are using. Others like geometrically simple low-complexity drawings of various objects [19].

Note that this automatically takes into account the context-dependency of regularity and beauty. Elements of many artfully designed Japanese gardens reflect environmental features such as silhouettes of surrounding mountains. The shared regularities make the description of both garden and environment shorter than the description of the garden by itself, concatenated with the description of the environment by itself. For similar reasons a mountain village of similar shacks made from nearby mountain rocks looks “nicer” in the eyes of many observers than a less homogenous, less compressible collection of houses with arbitrary, partially incompatible styles.

What’s beautiful is not necessarily interesting or aesthetically rewarding though. In fact, many artists agree that good art does not have to be beautiful. A beautiful thing is interesting only as long as it is new, that is, as long as the algorithmic regularity that makes it simple has not yet been fully assimilated by the adaptive observer who is still learning to compress the data better. This is the subject of the next section.

1.2 Novel Patterns and Subjective Interestingness

Babies, scientists, artists, and others do not use the same compression method forever, but improve it all the time, continually exploiting newly discovered regularities in the incoming data stream. It is possible to build *artificial* self-improving compression algorithms based on artificial neural networks and other machine learning techniques [17, 16, 30, 20, 21, 24, 26, 25]. For both artificial and biological observers we may define the time-dependent subjective *Interestingness* or *Novelty* or *Surprise* or *Aesthetic Reward* or *Aesthetic Value* $I(D, O(t))$ of data D relative to observer O at time t by

$$I(D, O(t)) \sim \frac{\partial B(D, O(t))}{\partial t}, \quad (1)$$

the *first derivative* of subjective beauty: as the learning agent improves its compression algorithm, formerly apparently random data parts become subjectively more regular and beautiful, requiring fewer and fewer bits for their encoding. As long as this process is not over the data remains interesting, but eventually it becomes boring even if it is beautiful.

Discrete time variants of this concept have been implemented and / or described in detail earlier [17, 16, 30, 20, 21, 24, 26, 25]. One of the simplest ways of calculating subjective interestingness is this: count how many bits are needed to encode the data before and after learning; the difference (the number of *saved* bits) corresponds to the internal joy or intrinsic reward for having found a new, previously unknown regularity—a novel pattern.

1.3 Maximizing the Joy of Finding Novel Patterns

Everybody actively chooses the actions that influence the observed data: a scientist chooses his experiments, a baby its toys, a visual artist his colors, a dancer his moves, or any attentive system its next sensory input. Above we have identified three essential ingredients of a creative (possibly artificial) scientist or artist: (1) a predictor or compressor of the continually growing data history (exploiting currently *known* patterns), (2) a learning algorithm that continually improves the predictor or compressor (detecting *novel* patterns), (3) a formal measure of the compressor's progress (due to the learning algorithm), to calculate intrinsic or aesthetic rewards. What is still missing is (4) a reward optimizer translating those rewards into action sequences expected to maximize cumulative future reward, by *making* novel patterns (possibly through exploitation of processes in the environment). The field of Reinforcement Learning (RL) offers many (more or less powerful) methods for doing so [9, 8, 27]. Some of them were used in our earlier implementations and descriptions of curious, creative systems [17, 16, 30, 20, 21, 24, 26, 25].

1.4 Relation to External Reward

Of course, the real goal of many cognitive systems is not just to find novel patterns, but to solve externally given problems. Any formalizable problem can be phrased as a Reinforcement Learning (RL) problem for an agent living in a possibly unknown environment, trying to maximize the future external reward expected until the end of its possibly finite lifetime. The new millennium brought a few extremely general, even universal RL algorithms or universal problem solvers or universal Artificial Intelligences (AIs) that are optimal in various theoretical but not necessarily practical senses, e. g., [8, 27]. To the extent that search for novel patterns and compression progress as above are helpful, these universal methods will automatically discover and exploit such concepts. Then why bother at all writing down an explicit framework for creativity and active pattern search and curiosity-based experimentation?

One answer is that the present universal approaches sweep under the carpet certain problem-independent constant slowdowns, by burying them in the asymptotic notation of theoretical computer science. They leave open an essential remaining question: If the agent can execute only a fixed number of computational instructions per unit time interval (say, 10 trillion elementary operations per second), what is the best way of using them to get as close as possible to the recent theoretical limits of universal AIs, especially when external rewards are very rare, as is the case in many realistic environments? The premise of this paper is that the curiosity drive or "novel pattern search drive" is such a general and generally useful concept for limited-resource RL in rare-reward environments that it should be prewired, as opposed to be learnt from scratch, to save on (constant but possibly still huge) computation time. An inherent assumption of this approach is that in realistic worlds a better explanation of the past can only help to better predict the future, and to accelerate the search for solutions to externally given tasks, ignoring the possibility that curiosity may actually be harmful and "kill the cat."

Many of our creative RL machines simply add external and internal rewards, maxi-

mizing the expected future cumulative sum of both [16, 30, 21, 24, 26, 25]. In absence of external reward, or when there is no known way to further increase the expected external reward, such a machine essentially tries to maximize *true novelty* or *interestingness* or *aesthetic reward* or *artistic value*, the *first derivative* of subjective beauty or compressibility, the steepness of the learning curve. It will do its best to select action sequences expected to create observations yielding maximal expected future compression *progress*, given the limitations of both the compressor or predictor and the compressor improvement algorithm. It will learn to focus its attention and its actively chosen experiments on things that currently are not optimally compressible but are expected to become more compressible / predictable through additional learning. It will get bored by things that already are subjectively highly compressible (patterns that are not novel any more). It will also get bored by things that are currently incompressible but will apparently remain so, given the experience so far, or where the costs of making them subjectively compressible exceed those of making other things compressible, etc.

Note in this context that our above concepts of beauty and interestingness are limited and *pristine* in the sense that they are *not a priori* related to pleasure derived from external rewards or punishments. For example, some might claim that a hot bath on a cold day feels “beautiful” due to rewards for achieving prewired target values of external temperature sensors (external in the sense of: outside the brain which is controlling the actions of its external body). Or a song may be called “beautiful” for emotional reasons by some who associate it with memories of external pleasure through their first kiss. Obviously this is not what we have in mind here—we are focusing on rewards of the intrinsic type based on learning progress, that is, the discovery of novel patterns.

2 Consequences of Novel Pattern Search

Let us discuss how many essential ingredients of intelligence and cognition can be viewed as natural by-products of the principles above.

2.1 Art & Music as By-Products of Novel Pattern Search

Works of art and music may have important purposes beyond their social aspects despite of those who classify art as superfluous. Good observer-dependent art deepens the observer’s insights about this world or possible worlds, unveiling previously unknown regularities in compressible data, connecting previously disconnected patterns in an initially surprising way that makes the combination of these patterns subjectively more compressible (art as an eye-opener), and eventually becomes known and less interesting. I postulate that the active creation and attentive perception of all kinds of artwork are just by-products of our principle of interestingness and curiosity yielding reward for compressor improvements. Let us elaborate on this idea in more detail, following the discussion in [21, 24, 26, 25].

Artificial or human observers must perceive art sequentially, and typically also actively, e.g., through a sequence of attention-shifting eye saccades or camera movements scanning a sculpture, or internal shifts of attention that filter and emphasize sounds made by a pianist, while suppressing background noise. Undoubtedly many

derive pleasure and rewards from perceiving works of art, such as certain paintings, or songs. But different subjective observers with different sensory apparatus and compressor improvement algorithms will prefer different input sequences. Hence any objective theory of what is good art must take the subjective observer as a parameter, to answer questions such as: Which sequences of actions and resulting shifts of attention should he execute to maximize his pleasure? According to our principle he should select one that maximizes the quickly learnable compressibility that is new, relative to his current knowledge and his (usually limited) way of incorporating / learning / compressing new data.

For example, which song should some human observer select next? Not the one he just heard ten times in a row. It became too predictable in the process. But also not the new weird one with the completely unfamiliar rhythm and tonality. It seems too irregular and contain too much arbitrariness and subjective noise. He should try a song that is unfamiliar enough to contain somewhat unexpected harmonies or melodies or beats etc., but familiar enough to allow for quickly recognizing the presence of a new learnable regularity or compressibility in the sound stream: A novel pattern! Sure, this song will get boring over time, but not yet.

The observer dependence is illustrated by the fact that Schönberg's twelve tone music is less popular than certain pop music tunes, presumably because its algorithmic structure is less obvious to many human observers as it is based on more complicated harmonies. For example, frequency ratios of successive notes in twelve tone music often cannot be expressed as fractions of very small integers. Those with a prior education about the basic concepts and objectives and constraints of twelve tone music, however, tend to appreciate Schönberg more than those without such an education.

All of this perfectly fits our principle: The current compressor of a given subjective observer tries to compress his history of acoustic and other inputs where possible. The action selector tries to find history-influencing actions that improve the compressor's performance on the selected historic data. The interesting musical and other subsequences are those with previously unknown yet learnable types of regularities, because they lead to compressor improvements. The boring patterns are those that are already perfectly known or arbitrary or random, or whose structure seems too hard to understand.

Similar statements not only hold for other dynamic art including film and dance (taking into account the compressibility of controller actions), but also for painting and sculpture, which cause dynamic pattern sequences due to attention-shifting actions of the observer.

2.2 No Objective “Ideal Ratio” Between Expected and Unexpected

Some of the previous attempts at explaining aesthetic experience in the context of information theory [2, 13, 1, 4, 14, 6] emphasized the idea of an “*ideal*” ratio between expected and unexpected information conveyed by some aesthetic object (its “*order*” vs its “*complexity*”). Note that our alternative approach does not have to postulate an objective ideal ratio of this kind. Instead, and unlike some of the previous works that already emphasized the significance of the subjective observer [4, 6, 5], our dynamic measure of interestingness reflects the *change* in the number of bits required to encode

an object, and explicitly takes into account the subjective observer’s prior knowledge as well as its limited compression *improvement* algorithm.

Hence the value of an aesthetic experience is not defined by the observed object *per se*, but by the algorithmic compression *progress* of the subjective, learning observer. Why didn’t early pioneers of aesthetic information theory put forward similar views? Perhaps because back then the fields of algorithmic information theory and adaptive compression through machine learning were still in their infancy?

2.3 Difference to Traditional Information Theory

Consider two extreme examples of uninteresting, unsurprising, boring data: A vision-based agent that always stays in the dark will experience an extremely compressible, soon totally predictable history of unchanging visual inputs. In front of a screen full of white noise conveying a lot of information and “novelty” and “surprise” in the traditional sense of Boltzmann and Shannon [28], however, it will experience highly unpredictable and fundamentally incompressible data. In both cases the data is boring [20, 24] as it does not allow for further compression progress—there is no novel pattern. Therefore we reject the traditional notion of surprise. Neither the arbitrary nor the fully predictable is *truly* novel or surprising—only data with still *unknown* algorithmic regularities are [17, 16, 30, 20, 21, 24, 26, 25]!

2.4 How Active Artists and Passive Observers of Art are Alike

Just as observers get intrinsic rewards for sequentially focusing attention on artwork that exhibits new, previously unknown regularities, the *creative* artists get reward for making it. For example, I found it extremely rewarding to discover (after hundreds of frustrating failed attempts) the simple geometric regularities that permitted the construction of the drawings in Figures 1 and 2. The distinction between artists and observers is blurred though. Both execute action sequences to exhibit new types of compressibility. The intrinsic motivations of both are fully compatible with our simple principle.

Some artists, of course, crave *external* reward from other observers, in form of praise, money, or both, in addition to the *intrinsic* compression improvement-based reward that comes from creating a truly novel work of art. Our principle, however, conceptually separates these two reward types, although our reinforcement learner can simply add them together, maximizing the expected future cumulative sum of both [16, 30, 20, 21, 24, 26, 25].

2.5 Science as By-Product of Novel Pattern Search

If the history of the entire universe were computable, and there is no evidence against this possibility [22], then its simplest explanation would be the shortest program that computes it. Unfortunately there is no general way of finding the shortest program computing any given data [11]. Therefore physicists have traditionally proceeded incrementally, analyzing just a small aspect of the world at any given time, trying to find

simple laws that allow for describing their limited observations better than the best previously known law, essentially trying to find a program that compresses the observed data better than the best previously known program. An unusually large compression breakthrough deserves the name *discovery*. For example, Newton's law of gravity can be formulated as a short piece of code which allows for substantially compressing many observation sequences involving falling apples and other objects. Although its predictive power is limited—for example, it does not explain quantum fluctuations of apple atoms—it still allows for greatly reducing the number of bits required to encode the data stream, by assigning short codes to events that are predictable with high probability [7] under the assumption that the law holds. Einstein's general relativity theory yields additional compression progress as it compactly explains many previously unexplained deviations from Newton's predictions. Most physicists believe there is still room for further advances, and this is what's driving their ongoing research [26, 25].

2.6 How Artists and Scientists are Alike

From our perspective, scientists are very much like artists. They actively select experiments in search for simple but new laws compressing the resulting observation history. In particular, the *creativity* of painters, dancers, musicians, pure mathematicians, physicists, can be viewed as a mere by-product of our exploration framework based on the compression progress drive. All of them try to create novel, non-random, non-arbitrary patterns with surprising, previously unknown regularities. For example, many physicists invent experiments to create data governed by previously unknown laws allowing to further compress the data. On the other hand, many artists combine well-known objects in a subjectively novel way such that the observer's subjective description of the result is shorter than the sum of the lengths of the descriptions of the parts, due to some previously unnoticed regularity shared by the parts.

What is the main difference between science and art? The essence of science is to *formally nail down* the nature of compression progress achieved through the discovery of a new regularity. For example, the law of gravity can be described by just a few symbols. In the fine arts, however, compression progress achieved by observing an artwork combining previously disconnected things in a new way (art as an eye-opener) may be *subconscious* and not at all formally describable by the observer, who may *feel* the progress in terms of intrinsic reward without being able to say exactly which of his memories became more subjectively compressible in the process.

The formal framework of previous papers [17, 16, 30, 20, 21, 24, 26, 25] is sufficiently formal to allow for implementation of our principle on computers. The resulting artificial observers will vary in terms of the computational power of their history compressors and learning algorithms. This will influence what is good art / science to them, and what they find interesting.

2.7 Jokes and Comedy and Other Sources of Fun

Just like other entertainers and artists, comedians also tend to combine well-known concepts in a novel way such that the observer's subjective description of the result

is shorter than the sum of the lengths of the descriptions of the parts, due to some previously unnoticed regularity shared by the parts.

In many ways the laughs provoked by witty jokes are similar to those provoked by the acquisition of new skills through both babies and adults. Past the age of 25 I learnt to juggle three balls. It was not a sudden process but an incremental and rewarding one: in the beginning I managed to juggle them for maybe one second before they fell down, then two seconds, four seconds, etc., until I was able to do it right. Watching myself in the mirror (as recommended by juggling teachers) I noticed an idiotic grin across my face whenever I made progress. Later my little daughter grinned just like that when she was able to stand on her own feet for the first time. All of this makes perfect sense within our algorithmic framework: such grins presumably are triggered by intrinsic reward for generating a data stream with previously unknown novel patterns, such as the sensory input sequence corresponding to observing oneself juggling, which may be quite different from the more familiar experience of observing somebody else juggling, and therefore truly novel and intrinsically rewarding, until the adaptive predictor / compressor gets used to it.

3 Concrete Implementations of Systems that Find or Create Novel Patterns

Predictors and compressors are closely related. Any type of partial predictability of the incoming sensory data stream can be exploited to improve the compressibility of the whole. Therefore the systems described in the first publications on artificial curiosity and creativity [17, 16, 30] already can be viewed as examples of implementations of a compression progress drive that encourages the discovery or creation of novel patterns.

3.1 Reward for Prediction Error

Early work [17] described a predictor based on a recurrent neural network (in principle a rather powerful computational device, even by today's machine learning standards), predicting sensory inputs including reward signals from the entire history of previous inputs and actions. The curiosity rewards were proportional to the predictor errors, that is, it was implicitly and optimistically assumed that the predictor will indeed improve whenever the controller learns to go to places where the prediction error is high.

3.2 Reward for Predictor Improvements

Follow-up work [16] pointed out that this approach may be inappropriate, especially in probabilistic environments: one should not focus on the errors of the predictor, but on its improvements. Otherwise the system will concentrate its search on those parts of the environment where it can always get high prediction errors due to noise or randomness, or due to computational limitations of the predictor, which will prevent improvements of the subjective compressibility of the data. While the neural predictor of the implementation described in the follow-up work was indeed computationally less powerful than the previous one [17], there was a novelty, namely, an explicit (neural) adaptive

model of the predictor’s improvements. This model essentially learned to predict the predictor’s changes. For example, although noise was unpredictable and led to wildly varying target signals for the predictor, in the long run these signals did not change the adaptive predictor parameters much, and the predictor of predictor changes was able to learn this. A standard RL algorithm [9] was fed with curiosity reward signals proportional to the expected long-term predictor changes, and thus tried to maximize information gain [3] by making novel patterns within the given limitations. In fact, we may say that the system tried to maximize an approximation of the (discounted) sum of the expected first derivatives of the data’s subjective predictability, thus also maximizing an approximation of the (discounted) sum of the expected changes of the data’s subjective compressibility.

3.3 Relative Entropy between Agent’s Prior and Posterior

Additional follow-up work yielded an information theory-oriented variant of the approach in non-deterministic worlds [30] (1995). The curiosity reward was again proportional to the predictor’s surprise / information gain, this time measured as the Kullback-Leibler distance [10] between the learning predictor’s subjective probability distributions before and after new observations - the relative entropy between its prior and posterior. In 2005 it was demonstrated experimentally that our approach explains certain patterns of human visual attention better than certain previous approaches.

Note that the concepts of Huffman coding [7] and relative entropy between prior and posterior immediately translate into a measure of learning progress reflecting the number of saved bits—a measure of improved data compression.

Note also, however, that the naive probabilistic approach to data compression is unable to discover more general types of *algorithmic* compressibility [11]. For example, the decimal expansion of π looks random and incompressible but isn’t: there is a very short algorithm computing all of π , yet any finite sequence of digits will occur in π ’s expansion as frequently as expected if π were truly random, that is, no simple statistical learner will outperform random guessing at predicting the next digit from a limited time window of previous digits. More general *program* search techniques are necessary to extract the underlying algorithmic regularity.

3.4 Learning Programs & Skills Through Zero Sum Reward Games

More recent work [20] (1997-2002) greatly increased the computational power of controller and predictor by implementing them as co-evolving, symmetric, opposing modules consisting of self-modifying probabilistic programs written in a universal programming language. The internal storage for temporary computational results of the programs was viewed as part of the changing environment. Each module could suggest experiments in the form of probabilistic algorithms to be executed, and make confident predictions about their effects by betting on their outcomes, where the ‘*betting money*’ essentially played the role of the intrinsic reward. The opposing module could reject or accept the bet in a zero-sum game by making a contrary prediction. In case of acceptance, the winner was determined by executing the algorithmic experiment and checking its outcome; the money was eventually transferred from the surprised loser to

the confirmed winner. Both modules tried to maximize their money using a rather general RL algorithm designed for complex stochastic policies (alternative RL algorithms could be plugged in as well). Thus both modules were motivated to discover *truly novel* algorithmic patterns / compressibility, where the subjective baseline for novelty was given by what the opponent already knew about the world's repetitive patterns. The system learnt a complex, partly hierarchical structure of more and more complex skills or programs (solving self-generated tasks) on top of previously acquired simpler skills.

3.5 Optimal Systems Driven by Prediction Progress & Compression Progress

Several new variants of the basic principle of novel pattern creation were introduced more recently [21], in particular, theoretically optimal ways of measuring prediction progress / compression progress / surprise / novelty, as well as mathematically optimal variants of selecting action sequences or experiments within the corresponding framework of artificial curiosity and creativity [21, 24, 26, 25].

3.6 No Essential Difference to Human Artists & Scientists?

Our previous computer programs already incorporated (approximations of) the basic creativity principle. But do they really deserve to be viewed as rudimentary artists and scientists? The patterns they create are novel with respect to their own limited predictors and prior knowledge, but not necessarily relative to the knowledge of sophisticated adults. The main difference to human artists / scientists, however, may be only quantitative by nature, not qualitative:

1. The unknown learning algorithms of human predictors / compressors are presumably still better suited to real world data. Note, however, that there already exist *universal*, mathematically optimal (but not necessarily practically feasible) prediction and compression algorithms [8, 27], and that ongoing research is continually producing better and better *practical* prediction & compression methods, to be plugged into our creativity framework.
2. Humans may have superior reinforcement learning algorithms for maximizing rewards generated through compression improvements achieved by their predictors. Note, however, that there already exist *universal*, mathematically *optimal* (but not necessarily practically feasible) reward optimizing algorithms [8, 27], and that ongoing research is continually producing better and better *practical* reinforcement learning methods, to be plugged into our creativity principle.
3. Renowned human artists and scientists have had decades of training experiences involving a multitude of high-dimensional sensory inputs and motoric outputs, while our systems so far only had a few hours with very low-dimensional experiences in limited artificial worlds. This quantitative gap, however, will narrow as we are scaling up our systems.

4. Human brains still have vastly more raw computational power and storage capacity than the best artificial computers. Note, however, that this statement is unlikely to remain true for more than a few decades - currently each decade brings a hardware speed-up factor of roughly 1000.

3.7 Improving Real Reward Intake

Our references above demonstrated experimentally that the presence of intrinsic reward or curiosity reward actually can speed up the collection of *external* reward.

3.8 Other Implementations

Recently several researchers also implemented variants or approximations of the curiosity framework. Singh and Barto and coworkers focused on implementations within the option framework of RL [29], directly using prediction errors as curiosity rewards [17]. Additional implementations were presented at the 2005 AAAI Spring Symposium on Developmental Robotics.

4 Illustrations of Subjective Beauty and its *First Derivative* Interestingness

As mentioned above (Section 3.3), the probabilistic variant of our theory [30] (1995) was able to explain certain shifts of human visual attention. But we can also apply our approach to the complementary problem of *constructing* images that contain quickly learnable regularities, arguing again that there is no fundamental difference between the motivation of creative artists and passive observers of visual art (Section 2.4). Both create action sequences yielding interesting inputs, where interestingness is a measure of learning progress, for example, based on the relative entropy between prior and posterior (Section 3.3), or the saved number of bits needed to encode the data, or something similar (Section 3).

Here we provide examples of subjective beauty tailored to human observers, and illustrate the learning process leading from less to more subjective beauty. Due to the nature of the present written medium, we have to use visual examples instead of acoustic or tactile ones. Our examples are intended to support the hypothesis that unsupervised *attention* and the *creativity* of artists, dancers, musicians, pure mathematicians is just a by-product of our drive to find novel patterns and enable compression progress.

4.1 A Pretty Simple Face with a Short Algorithmic Description

Figure 1 depicts the construction plan of a female face considered '*beautiful*' by some human observers. It also shows that the essential features of this face follow a very simple geometrical pattern [24, 26, 25] that can be specified by very few bits of information. That is, the data stream generated by observing the image (say, through a sequence of eye saccades) is more compressible than it would be in the absence of such regularities. Although few people are able to immediately see how the drawing was

made in absence of its superimposed grid-based explanation, most do notice that the facial features somehow fit together and exhibit some sort of regularity. According to our postulate, the observer’s reward is generated by the conscious or subconscious discovery of this compressibility. The face remains interesting until its observation does not reveal any additional previously unknown regularities. Then it becomes boring even in the eyes of those who think it is beautiful—as has been pointed out repeatedly above, beauty and interestingness are two different things.

4.2 Another Drawing That Can Be Encoded By Very Few Bits

Figure 2 provides another example: a butterfly and a vase with a flower. It can be specified by very few bits of information as it can be constructed through a very simple procedure or algorithm based on fractal circle patterns [19]—see Figure 3. People who understand this algorithm tend to appreciate the drawing more than those who do not. They realize how simple it is. This is not an immediate, all-or-nothing, binary process though. Since the typical human visual system has a lot of experience with circles, most people quickly notice that the curves somehow fit together in a regular way. But few are able to immediately state the precise geometric principles underlying the drawing [21]. This pattern, however, is learnable from Figure 3. The conscious or subconscious discovery process leading from a longer to a shorter description of the data, or from less to more compression, or from less to more subjectively perceived beauty, yields intrinsic reward or internal joy or aesthetic value depending on the first derivative of subjective beauty, that is, the steepness of the learning curve.

5 Conclusion & Outlook

We pointed out that a surprisingly simple algorithmic principle for novel pattern search based on the notions of data compression and data compression *progress* informally explains fundamental aspects of attention, novelty, surprise, interestingness, curiosity, creativity, subjective beauty, jokes, and science & art in general. The crucial ingredients of the corresponding *formal* framework are (1) a predictor or compressor of the continually growing data history (exploiting already *known* patterns), (2) a learning algorithm that continually improves the predictor or compressor (detecting *novel* patterns that subsequently become *known* patterns), (3) a formal measure of the compressor’s progress (due to the learning algorithm), to calculate intrinsic or aesthetic rewards, (4) a reward optimizer translating those rewards into action sequences expected to maximize cumulative future reward, by *creating* novel patterns (possibly with help from the environment). To improve our previous implementations of these ingredients (Section 3), we will (A) study better adaptive compressors, in particular, recent, novel artificial recurrent neural networks and other general but practically feasible methods for making predictions; (B) investigate under which conditions learning progress measures can be computed both accurately and efficiently, without frequent expensive compressor performance evaluations on the entire history so far; (C) study the applicability of recent improved RL techniques in the fields of policy gradients, artificial evolution, and others.

5.1 Acknowledgments

Thanks to Herbert W. Franke, Marcus Hutter, Andy Barto, Jonathan Lansey, Julian Togelius, Faustino J. Gomez, Giovanni Pezzulo, Gianluca Baldassarre, Martin Butz, for useful comments that helped to improve this paper, or earlier papers on this subject.

References

- [1] M. Bense. *Einführung in die informationstheoretische Ästhetik. Grundlegung und Anwendung in der Texttheorie (Introduction to information-theoretical aesthetics. Foundation and application to text theory)*. Rowohlt Taschenbuch Verlag, 1969.
- [2] G. D. Birkhoff. *Aesthetic Measure*. Harvard University Press, 1933.
- [3] V. V. Fedorov. *Theory of optimal experiments*. Academic Press, 1972.
- [4] H. G. Frank. *Kybernetische Analysen subjektiver Sachverhalte*. Verlag Schnelle, Quickborn, 1964.
- [5] H. G. Frank and H. W. Franke. *Ästhetische Information. Estetika informacio. Eine Einführung in die kybernetische Ästhetik*. Kopäd Verlag, 2002.
- [6] H. W. Franke. *Kybernetische Ästhetik. Phänomen Kunst (3rd edition)*. Ernst Reinhardt Verlag, Munich, 1979.
- [7] D. A. Huffman. A method for construction of minimum-redundancy codes. *Proceedings IRE*, 40:1098–1101, 1952.
- [8] M. Hutter. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*. Springer, Berlin, 2004. (On J. Schmidhuber’s SNF grant 20-61847).
- [9] L. P. Kaelbling, M. L. Littman, and A. W. Moore. Reinforcement learning: a survey. *Journal of AI research*, 4:237–285, 1996.
- [10] S. Kullback. *Statistics and Information Theory*. J. Wiley and Sons, New York, 1959.
- [11] M. Li and P. M. B. Vitányi. *An Introduction to Kolmogorov Complexity and its Applications (2nd edition)*. Springer, 1997.
- [12] B. Mandelbrot. *The Fractal Geometry of Nature*. Freeman and Co., San Francisco, 1982.
- [13] A. Moles. *Information Theory and Esthetic Perception*. Univ. of Illinois Press, 1968.
- [14] F. Nake. *Ästhetik als Informationsverarbeitung*. Springer, 1974.
- [15] K. R. Popper. *All Life Is Problem Solving*. Routledge, London, 1999.

- [16] J. Schmidhuber. Curious model-building control systems. In *Proceedings of the International Joint Conference on Neural Networks, Singapore*, volume 2, pages 1458–1463. IEEE press, 1991.
- [17] J. Schmidhuber. A possibility for implementing curiosity and boredom in model-building neural controllers. In J. A. Meyer and S. W. Wilson, editors, *Proc. of the International Conference on Simulation of Adaptive Behavior: From Animals to Animats*, pages 222–227. MIT Press/Bradford Books, 1991.
- [18] J. Schmidhuber. *Femmes fractales*, 1997.
- [19] J. Schmidhuber. Low-complexity art. *Leonardo, Journal of the International Society for the Arts, Sciences, and Technology*, 30(2):97–103, 1997.
- [20] J. Schmidhuber. Exploring the predictable. In A. Ghosh and S. Tsutsui, editors, *Advances in Evolutionary Computing*, pages 579–612. Springer, 2002.
- [21] J. Schmidhuber. Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. *Connection Science*, 18(2):173–187, 2006.
- [22] J. Schmidhuber. Randomness in physics. *Nature*, 439(3):392, 2006. Correspondence.
- [23] J. Schmidhuber. Alle berechenbaren Universen (All computable universes). *Spektrum der Wissenschaft Spezial (German edition of Scientific American)*, (3):75–79, 2007.
- [24] J. Schmidhuber. Simple algorithmic principles of discovery, subjective beauty, selective attention, curiosity & creativity. In *Proc. 10th Intl. Conf. on Discovery Science (DS 2007), LNAI 4755*, pages 26–38. Springer, 2007. Joint invited lecture for *ALT 2007 and DS 2007*, Sendai, Japan, 2007.
- [25] J. Schmidhuber. Driven by compression progress: A simple principle explains essential aspects of subjective beauty, novelty, surprise, interestingness, attention, curiosity, creativity, art, science, music, jokes. In G. Pezzulo, M. V. Butz, O. Sigaud, and G. Baldassarre, editors, *Anticipatory Behavior in Adaptive Learning Systems. From Psychological Theories to Artificial Cognitive Systems*, volume 5499 of *LNCS*, pages 48–76. Springer, 2009.
- [26] J. Schmidhuber. Simple algorithmic theory of subjective beauty, novelty, surprise, interestingness, attention, curiosity, creativity, art, science, music, jokes. *SICE Journal of the Society of Instrument and Control Engineers*, 48(1):21–32, 2009.
- [27] J. Schmidhuber. Ultimate cognition à la Gödel. *Cognitive Computation*, 1(2):177–193, 2009.
- [28] C. E. Shannon. A mathematical theory of communication (parts I and II). *Bell System Technical Journal*, XXVII:379–423, 1948.

- [29] S. Singh, A. G. Barto, and N. Chentanez. Intrinsically motivated reinforcement learning. In *Advances in Neural Information Processing Systems 17 (NIPS)*. MIT Press, Cambridge, MA, 2005.
- [30] J. Storck, S. Hochreiter, and J. Schmidhuber. Reinforcement driven information acquisition in non-deterministic environments. In *Proceedings of the International Conference on Artificial Neural Networks, Paris*, volume 2, pages 159–164. EC2 & Cie, 1995.

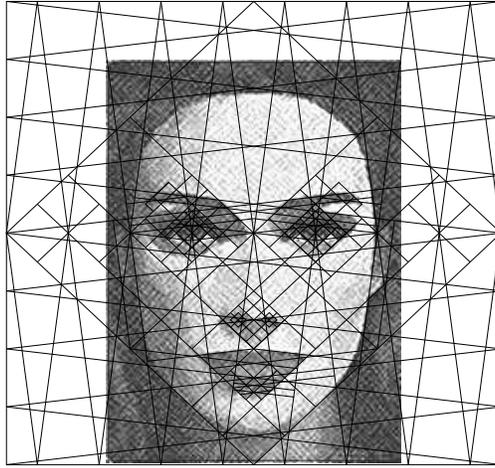


Figure 1: Previously published construction plan [24, 26, 25] of a female face (1998). Some human observers report they feel this face is 'beautiful.' Although the drawing has lots of noisy details (texture etc) without an obvious short description, positions and shapes of the basic facial features are compactly encodable through a very simple geometrical scheme, simpler and much more precise than ancient facial proportion studies by Leonardo da Vinci and Albrecht Dürer. Hence the image contains a highly compressible algorithmic regularity or pattern describable by few bits of information. An observer can perceive it through a sequence of attentive eye movements or saccades, and consciously or subconsciously discover the compressibility of the incoming data stream. How was the picture made? First the sides of a square were partitioned into 2^4 equal intervals. Certain interval boundaries were connected to obtain three rotated, superimposed grids based on lines with slopes ± 1 or $\pm 1/2^3$ or $\pm 2^3/1$. Higher-resolution details of the grids were obtained by iteratively selecting two previously generated, neighboring, parallel lines and inserting a new one equidistant to both. Finally the grids were vertically compressed by a factor of $1 - 2^{-4}$. The resulting lines and their intersections define essential boundaries and shapes of eyebrows, eyes, lid shades, mouth, nose, and facial frame in a simple way that is obvious from the construction plan. Although this plan is simple in hindsight, it was hard to find: hundreds of my previous attempts at discovering such precise matches between simple geometries and pretty faces failed.

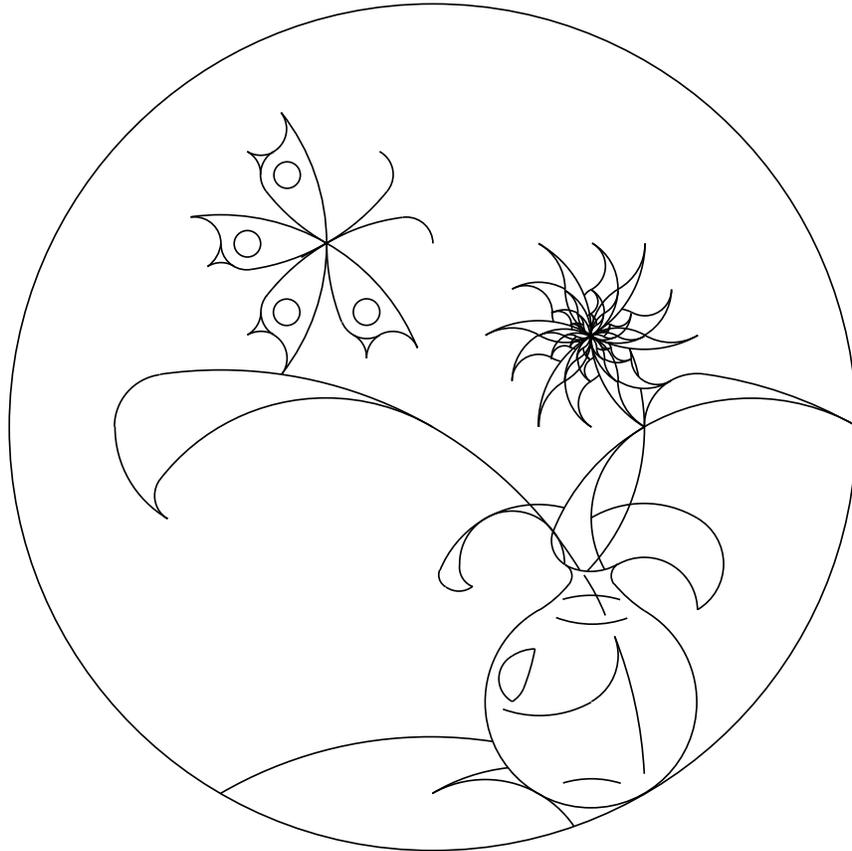


Figure 2: Image of a butterfly and a vase with a flower, reprinted from *Leonardo* [19, 21, 26, 25]. An explanation of how the image was constructed and why it has a very short description is given in Figure 3.

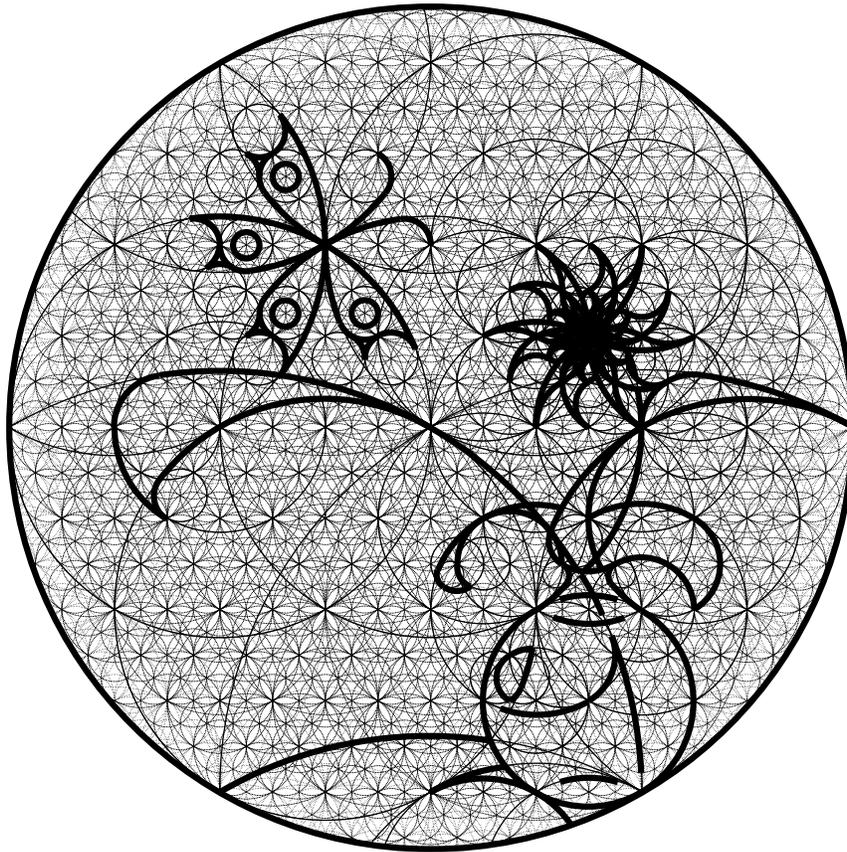


Figure 3: Explanation of how Figure 2 was constructed through a very simple algorithm exploiting fractal circles [19, 21, 26, 25]. The frame is a circle; its leftmost point is the center of another circle of the same size. Wherever two circles of equal size touch or intersect are centers of two more circles with equal and half size, respectively. Each line of the drawing is a segment of some circle, its endpoints are where circles touch or intersect. There are few big circles and many small ones. In general, the smaller a circle, the more bits are needed to specify it. The drawing is simple (compressible) as it is based on few, rather large circles. Many human observers report that they derive a certain amount of pleasure from discovering this simplicity. The observer's learning process causes a reduction of the subjective complexity of the data, yielding a temporarily high derivative of subjective beauty: a temporarily steep learning curve. (Again I needed a long time to discover a satisfactory and rewarding way of using fractal circles to create a reasonable drawing.)