Novelty Processing and Emotion: Conceptual Developments, Empirical Findings and Virtual Environments

Didier Grandjean and Christopher Peters

Abstract Novelty detection is a crucial ability of organisms to detect changes in the environment and to adapt their behaviours accordingly. In this chapter we review a conceptual framework of novelty detection informed by cognitive neuroscience and cognitive psychology. The relationship between attentional processes and novelty detection is also discussed and developed, supported by a case study highlighting methods for implementing a novelty detection capability for artificial agents in virtual environments.

1 Introduction

The growing knowledge about the functioning of human central nervous system and the possibilities for modelling such systems, at least in part, have created the potential for interdisciplinary work, for example, on the implementation of cognitive or emotional processes for artificial agents inhabiting virtual environments. The collaboration between cognitive neuroscience researchers and computer scientists working on virtual environments is important, supporting both the creation of realistic behaviours for agents inhabiting virtual environments and also allowing, through reconstruction attempts, testing of high-level behaviour implications of neuroscience results based on simple virtual analogies of real systems. In this chapter, we review the cognitive and cognitive neuroscience knowledge about novelty processing and provide an example of the first steps towards a computational implementation for supporting the dynamic interactions of virtual agents with their environment. Novelty processing is an essential ability of organisms to process novel information in order to increase their knowledge about their own environment. This mechanism allows organisms to detect changes in the environment and, with

D. Grandjean (✉)
Neuroscience of Emotion and Affective Dynamics Laboratory, Department of Psychology, University of Geneva, Geneva, Switzerland
e-mail: Didier.Grandjean@unige.ch

P. Petta et al. (eds.), *Emotion-Oriented Systems*, Cognitive Technologies, 441
DOI 10.1007/978-3-642-15184-2_23,

memory processes, to learn invariants in the environment. Even at a low level and for basic living systems, like nematodes, novelty processing is important, allowing these organisms to detect changes in their environment in order to allocate resources (i.e. energy), to avoid potentially dangerous situations or to orient towards a specific location space (i.e., respectively, to avoid some increase of salt gradient or to orient towards a source of food).

2 Novelty and Emotional Processes

The importance of novelty in emotional processes was suggested by appraisal theorists in the 1980s (Lazarus, 1991; Scherer, 1984). Novelty is not only closely related to surprise (which could be either positive or negative) but could be also determinant, at early stages of processing, in appraisal processes for several other emotions. In the context of emotion theories, the component process model (CPM) proposes emotional processes as massive synchronised changes occurring at the same time in the different sub-systems of the organism (Grandjean et al., 2008; Scherer, 2001). These systems have been described as the cognitive, peripheral, motivational, expressive and the monitor sub-systems interacting dynamically. In this chapter we will focus on a specific process, the detection of novelty, embedded into the cognitive system in the emotional process. In the context of appraisal processes, novelty detection has been proposed as the first subprocess occurring in time when an organism appraises its environment. Actually, Scherer has formulated the notion of a fixed sequence of different appraisals in a theoretical proposition in 1984 (Scherer, 1984). Based on phylogenetical, ontogenical and logic arguments (Sander et al., 2005) the main idea is that the different appraisal subprocesses have to be organised in a fixed sequence in time. Every check would compute a preliminary closure impacting on the different other components. This sequential view is compatible with further parallel processing, for example, to represent novelty at higher levels. Moreover, we have developed a series of arguments to the sequential nature of appraisal process based on cognitive neuroscience empirical studies (see also Grandjean, 2005). Recent findings based on electroencephalography (EEG) have supported this view (Grandjean and Scherer, 2008). In two EEG experiments in the visual modality, different appraisals were manipulated in order to test the sequential hypothesis; novelty was manipulated by the probability of the occurrence of stimuli, intrinsic pleasantness was operationalised by positive, negative and neutral pictures (IAPS pictures;(Lang et al., 2005)), goal relevance was manipulated by the participant's task and finally goal conduciveness was manipulated by reward–punishment contingencies of different types of visual stimuli. Based on several types of EEG analysis, ERPs, topographical analysis and frequency analysis, the results indicate (1) a first modulation of brain electrical field at 50–100 ms related to novelty manipulation and (2) a later modulation of 150 ms related to intrinsic pleasantness and later components related to goal relevance and goal conduciveness (Grandjean and Scherer, 2008).

In the following paragraphs, we present and discuss a new typology of the concept of novelty. Indeed, in the literature in cognitive psychology and cognitive neuroscience, this term has been used in our point of view to refer to many different kinds of phenomena. For the purposes of clarification, we propose the following distinctions between different types of novelty processes. We distinguish between four major kinds of novelty processes, namely perceptual novelty, partial novelty, contextual novelty and semantic novelty. We define perceptual novelty as the construction of a new representation of an object never perceived in the past by the organism and requiring a new encoding in short-term and long-term memory. This "real" novelty detection has been largely neglected in cognitive and neuro-cognitive research. Indeed, EEG measures, often used to investigate novelty processes, require a number of repetitions of different stimuli types in order to compute event-related potentials (ERPs) and perform the subsequent statistical analysis. The investigation of perceptual novelty is not easy due to the difficulty of creating a series of totally new stimuli which are not different in terms of their basic physical features, such as luminance, visual complexity, content in spatial frequency; several confounds are possible and may be attributed to novelty effects while they may in fact be better related to basic physical differences between stimuli. For example, if someone wants to test the ERPs differences between familiarity and novelty, the different stimuli used in the experimental conditions have to be similar and controlled for basic physical visual features (Delplanque et al., 2007).

Partial novelty would be involved when an organism perceives an object that looks like a previously perceived object, but which has some perceptual differences for one or several characteristics. For example, if a baby used to play with a red ball, a new violet ball will not be totally new (in terms of perceptual novelty defined just above); it is only one or several of these characteristics which are new; the categorisation processes are very important in this concept of partial novelty. The fact that organisms are able to categorise different percepts in terms of one category, as, for example 'chairs', is a crucial ability allowing the extraction of the invariant in different dimensions; in this example, a functional utility, i.e. an object on which you can sit. This partial novelty would be sensitive to the boundaries of the category and then influenced by the mental setting of the individual, for example in terms of needs or desires.

Contextual novelty refers to situations in which an object is perceived in a new context or has emerged in a given stable context. For example, a ball on the table, when such objects are usually experienced on the floor, would induce a novelty detection phenomenon. Another example is the case of a familiar environment in which something new appears: on a monitor board a red light turns on to signal something unusual taking place; in this case the board is known, the red light as well, but the timing concomitance between the board and the red light is a new event in the present context. This type of novelty has been extensively investigated in cognitive neuroscience research, particularly using EEG. The need to have several dozen trials for one given experimental condition in EEG to compute ERPs has induced a bias in cognitive neuroscience towards the study of this type of novelty. In this kind of traditional EEG experiment, called the 'oddball paradigm', a flow of background stimuli (identical stimuli repeatedly presented to the participant) is interrupted by a 'new' stimulus, a deviant one. In the auditory domain, for example, the ERP component, called mismatch negativity (MMN), corresponding to a negative electrophysiological component appears when the deviant stimuli are presented (for example, a 500 Hz pure tone in a flow of 250 Hz tone Naantanen et al., 2007). Finally, semantic novelty refers to a situation in which the relationships between the objects or the concepts are organised in a new manner and have never been perceived as such in the past. The fact that individuals are able to create a new tool or a new concept from a series of well-known objects or ideas is characteristic of this type of novelty. Notice that this process is related to an active process involving an ability to create something new from something well known. The example of the use of stones as a tool to obtain food from different animals is exactly what we refer about this semantic novelty (Savage-Rumbaugh et al., 2003). This type of novelty is strongly related to creativity and the organisms abilities in terms of high cognitive computations.

Scherer has suggested distinguishing between different kinds of subprocesses related to novelty detection. Namely, he has proposed that suddenness, familiarity and predictability are particularly important in the processing of novelty detection (Scherer, 2001). Three different features of novelty have been proposed: (i) suddenness or abruptness of onset (Tulving and Kroll, 1995; Tulving et al., 1994), often coupled with high stimulation intensity, producing an orientation response (Siddle and Lipp, 1997); (ii) familiarity with the object or event (Habib et al., 2003), generally based on schema matching; and (iii) predictability, as based on past observations of regularities and probabilities for specific events. These three processes could be involved for each type of novelty that we described above. For example, suddenness is closely related to attentional processes and the orientation of attention and could be detected automatically or be processed at a higher level; for example, with a conscious perception of suddenness and reappraisal (see also below). The evaluation of familiarity and unpredictability is obviously involved for each type of novelty. These processes are strongly related to memory and could occur at different levels of processing: reflex, schematic or conceptual (Leventhal and Scherer, 1987). As already mentioned above, in the context of cognitive neuroscience, an important corpus of studies has investigated the modulations of different neural networks related to the experimental manipulation of novelty. The ERPs technique, based on the electrical signals of the brain, is one of the methods used to address the temporal dynamic of novelty processing at the brain level. Using this technique, cognitive researchers have addressed the brain mechanisms involved in novelty processes in contrast to the processes related to stimulus repetition. Stimuli repetition is classically accompanied by a decrease of neuronal activity in cortical and subcortical regions (Hensen and Rugg, 2003; Ranganath and Rainer, 2003; Ringo, 1996). A well-known electrophysiological component, the P300, has been extensively studied in the context of novelty processing (Friedman et al., 1993, 2001). This positive component, appearing at 300 ms after the onset of the visual stimuli, has been divided into two more specific electrophysiological components: the P3a and the P3b (Comerchero and Polich, 1999; Polich, 2007; Polich, 1988). Several studies have been demonstrated that the P3a is more prominent on the anterior part of the scalp compared to the

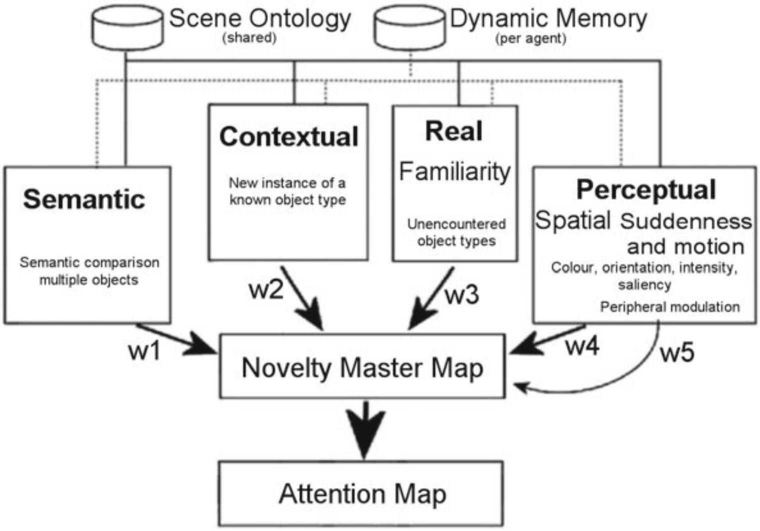

Fig. 1 Different types of novelty and their integration into a master map for guiding attention

P3b and that the former precedes in time the P3b. The shape of the P3a and its intensity are correlated to the novelty experimental manipulations while the P3b characteristics are more affected by the task performed by participants such as target detection task (Fig. 1; for a review see Ranganath and Rainer, 2003). The P3a and P3b neuronal generators are mainly located in the anterior and temporal regions. Human intracranial recordings have shown that anterior and temporal evoked potentials are compatible with the P3a or novelty P3 (Halgren et al., 1995a, b). More specifically, evoked responses have been found in the orbital, ventrolateral and dorsolateral prefrontal regions; similarly, different temporal regions have been shown being modulated by novelty, namely the medial part of the temporal region; perirhinal and posterior parahippocampal cortex, the hippocampus, the temporoparietal cortex and the cingulate gyrus. In the same way, clinical neuropsychology studies have confirmed these findings showing a decrease in novelty effects on the P3a for patients with lateral prefrontal (Knight, 1984), lateral temporoparietal or posterior medial temporal lobe lesions resulting from strokes (Knight, 1996). Functional magnetic resonance imaging (fMRI) studies have shown that the ventrolateral prefrontal cortex, cingulate gyrus and anterior insula are strongly involved in novelty processing (Knight and Nakada, 1998). Indeed, the distributed neuronal network involved in novelty processing includes areas in the lateral prefrontal cortex, orbital prefrontal, anterior insular and anterior temporal cortex, as well as temporoparietal cortex, medial temporal areas along the parahippocampal gyrus, hippocampal formation, amygdala and cingulated gyrus. A recent study investigating novelty and target processes using EEG and fMRI has shown that, despite a considerable overlap of regions activated during novelty and target processing, bilateral superior temporal and right inferior frontal areas are more activated in the novelty processing condition (Knight and Nakada, 1998).

Novelty detection induces an orientation response in the organism (Bernstein, 2002; Scherer, 2001); this orientation response is the behavioural part of the related attentional process effects and the subsequent consequences; that is, the recruitment of attentional and cognitive resources towards the new stimuli (Bernstein, 2002; Desire et al., 2004). The relationships between novelty detection and attentional processes have to be very intensive and modulated by the context of novelty appearance. Indeed, novelty detection may be related to the current goal needs of an individual or not. The modulation of the attentional processes should be different in these two extreme cases: when an organism detects something new in its environment independently and not related to these current task goals or needs, the organism has to inhibit this new information, especially if it is not relevant. In this case, new information is like a distractor and may induce inhibition processes in order to focus the organism's attention on the current task goals. The second case concerns new information appearance which could be related to the current task goals or needs of the individual; in this case the new information has to be processed deeply and involves sustained voluntary attentional processes. Indeed, several different kinds of attentional processes are recruited in novelty detection, dependent on the current task goals or needs of the individual. The unfolding of novelty processes can, indeed, include several steps which can occur at different levels of processing: (1) an early process involves the construction of a new representation; a first step of the encoding of a new stimuli; (2) the relatively automatic detection of unfamiliar, novel stimuli (indexed by the N2 for example) and the modulation of attentional processes; (3) the voluntary allocation of resources determined by the broader context in which a novel event occurs (indexed by the P3); and (4) the sustained processing of novelty, indexed by late-positive slow-wave activity (Chong et al., 2008). The first and second steps of information processing described above are mainly related to automatic detection and could take place at a low level of processing; especially at the reflexive level, eventually also at the schematic level (Leventhal and Scherer, 1987). The third and fourth steps would be more related to endogenous attentional processes involving the voluntary allocation of resources to process in greater depth the new information detected by the previous steps of information processing; see the schematic graph of this process.

As we have described above, novelty processing is not a unitary phenomenon, but may dissociate in different kinds of processes, not only at the conceptual level but also in terms of temporal dynamic processing and the relationship with attentional processes. At the conceptual level, we distinguish different types of novelty processing, allowing not only the implementation of different types of novelty in virtual environments, but also to clarify different concepts of novelty useful for experimental research (perceptual, contextual, partial and semantic novelty). In terms of timing, an understanding of the unfolding of novelty detection is important, allowing researcher to distinguish between the early stages of processing and later ones involving different kinds of attentional processes: these two steps of information processing are related to different levels of processing, with early stages related to exogenous attentional processes and later stages more related to endogenous attentional processes. Processing novelty in the environment is not independent of individual characteristics of the organism. Indeed, Cloninger et al. (1993) have proposed a psychobiological model including the concept of novelty-seeking defining the tendency for a given individual to orient and prefer novel stimuli or new situations. This temperament predisposition to search for and process preferentially novel information has been demonstrated to be related to different kinds of behaviours, like addiction and aggression (Ball, 2004). The predisposition of novelty seeking and other characteristics related to temperament and personality and their relationship with information processing should be taken into account in virtual environments to increase the richness of the possible behaviours that the avatars and agents can produce, not only in different situations but also in terms of individual differences. In the next section we describe how to apply the concepts defined above to present an outline for implementing novelty processing in virtual environments.

3 Novelty in Artificial Systems

The ability to detect novel elements of the environment is of prime importance for the survival and functioning of humans and other animals in dynamic and complex environments. As we have seen in Sect. 2, the growing knowledge about the functioning of the human central nervous system and the possibilities for the computational modelling of subsets of such systems have created the potential for interdisciplinary work, including computational implementations of cognitive and emotional processes. These implementations have been applied in many different domains, including data mining, surveillance, intrusion detection and robotics and have been based on a selection of approaches (Markou and Singh, 2009) and methodologies (Hodge and Austin, 2004). In the research literature, work relating to computational novelty detection has been referred to under a host of different names, including outlier detection and noise detection; see Hodge and Austin (2004, for overview.

An area of particular relevance to our discussion is robotics. Here, the concept of novelty has been linked to comparisons between pre-acquired memorised environment representations and current sensory data of robots in order to detect deviations from normal patterns. For example, Marsland et al. (2000) have presented an autonomous robot that senses an environment through sonar sensors and produces a novelty measure for each scan relative to the model it has learned. Metrics related to habituation are used to assign novelty values to perceptions. This concept is extended in Marsland et al. (2002), where different novelty filters are used in a variety of different contexts, allowing the context in which stimuli appear to be a factor in novelty judgments. More recently Neto and Nehmzow (2007) have combined an attention model with a novelty filter in order to process only salient locations, rather than the whole scene. This helps to improve the performance of such systems by allowing them to allocate extra processing resources to those stimuli deemed to be of higher potential relevance.

Novelty detection is particularly significant to computational agents if they are autonomous and are not granted full access to the scene database, but rather must process it through synthetic perception (see (Chapter “Fundamentals of Agent Perception and Attention Modelling”)). Instead of considering real systems, such as mobile robots, we consider virtual agents embedded inside virtual environments.

4 Case Study: Novelty Detection in Virtual Environments

In this case study, we describe one application of computational novelty in more detail: generating sensible attentive behaviours for autonomous agents in virtual environments. It provides an outline of how one might construct and integrate novelty detection for real-time virtual agents operating in virtual environments. These agents have a graphical 3D appearance and interact within the confines of a 3D virtual environment. They are often humanoid in appearance, in the case of the agents described here, have a degree of autonomy and are endowed with a synthetic perception for filtering and sensing their virtual surroundings. It should be noted that the nature of novelty processing as it relates to these agents is much divorced from the real systems described earlier in this chapter: in many ways, the real-time lightweight systems described here are very different and are still extremely inflexible when compared with their real-world counterparts and also more complex computational systems. Nonetheless, the aims and functional requirements are very similar for such agents, who are autonomous to some degree, do not have full access to the world database, inhabit complex virtual environments and attempt to operate under real-time constraints with very limited processing resources.

Here we list components for supporting novelty processing, a technical framework into which novelty processing can be placed and example applications. The overall purpose of the novelty system is to contribute to a prioritisation scheme for agents, allowing them to select and process only those details of potential relevance and importance to them according to the context of the situation or task: novelty is but one factor that may contribute to this – see (Chapter “Fundamentals of Agent Perception and Attention Modelling”) for possible others.

4.1 Supporting Capabilities

At least three fundamental capabilities are necessary for supporting novelty detection in computational systems:

1. An internal model of ‘normality’ based on past experience, e.g. from the external environment. There are numerous ways in which a model may be formed, for example as a static system with manually predefined details, or dynamic and learned based on exploratory behaviours, or blends of both.

2. A sensing system for sampling information from the environment. This may be based on different types of sensors, e.g. a simple temperature monitor, or more sophisticated camera-based sensor utilising computer vision techniques.
3. A comparison operation for comparing new incoming sensory data with the internal model of the environment, based on similarity (i.e. distance) metrics.

The existence of an internal model of normality which is subjective, based on experience and need not coincide with the real state of the environment, implied the need for some form of internal storage mechanism. The memory model (Sect. 4.1.1) fulfils this role. The state of the environment is sampled by a sensing (Sect. 4.1.2) system to provide a model of normality for the agent, against which newly encountered stimuli may be compared.

4.1.1 Memory Model

The memory system allows agents to store or create a model of the environment that can be regarded as 'normal' to them. The memory system described here (Peters, 2006) consists of two components – a static area and a dynamic object-based area. The static area of the agent's memory consists of an object ontology – it could be considered to be part of the agent's pre-defined long-term memory. Here, it is assumed that the ontology is static and no learning takes place, i.e. ontology node values do not change during the simulation, and the agent cannot add new object categories at runtime.

The dynamic area of the memory system provides a storage for object percepts that have been previously sensed by the agent, in addition to memory management descriptors. Unlike the real world, where defining the concept and nature of an object from input is a tenuous task, e.g. using vision approaches to segment a scene, in the virtual environment, the task is simplified as objects are defined by the creator of the scene when it is being constructed.

The role of the memory system and scene ontology is to store a history of familiarity with objects, such as how often they have been previously encountered and the degree to which they have been scrutinised in the dynamic area of memory, and the semantic similarities between objects in the static.

4.1.2 Synthetic Vision

A synthetic vision module samples stimuli from the environment in a snapshot manner by means of an orientable synthetic vision sensor that is locked to the gaze direction of the agent, to create an 'active vision' system – see Sect. 4.3. This sensor renders the scene from the viewpoint of the agent in two modes, allowing for two different basic types of scene representation to exist: spatial and object. By combining the full-coloured and false-coloured representations, the agent has access to a view-dependent representation of the scene, so it is capable of processing both object and spatial information from scenes (see (Chapter "Fundamentals of Agent Perception and Attention Modelling")).

4.2 Implementing Novelty Detection

The memory and synthetic perception capabilities make it possible to create a limited computational implementation of each of the four types of novelty previously described (see Fig. 1 and Peters and Grandjean 2008). As we have seen, defining novelty is a difficult prospect, as are the challenges awaiting those who attempt computational implementations. There are many ways in which items or events may appear novel, so an attempted categorisation is an imperative for supporting attempts towards the creation of computational novelty systems.

4.2.1 Real Novelty

Real novelty refers to a new object instance that has never been or rarely encountered before, although it may fall into a known category. For example, an elephant may be novel because one has never been observed before or to a lesser extent, a car may be deemed novel as this particular one (or instance) has not been observed before, despite the agent's familiarity with road vehicles.

In this case, the agent's dynamic memory system (see Sect. 4.1.1) acts as storage: An uncertainty metric is stored for each object and is a simplification used to represent the well formedness and accuracy of the agent's internal representation with respect to the actual state of the object being perceived. As the agent attends to an object, its uncertainty level is reduced in memory in order to signify the agent's increased knowledge about the object's state. Other information regarding the state of the object is also stored in memory – for example position and velocity. These allow the agent to have its own internal representation of the virtual environment that does not necessarily match its actual state. For example, if the agent is not attending to an object when its position changes, the object will maintain its old position (now incorrect) in the agent's memory.

4.2.2 Perceptual Novelty

Perceptual novelty operations are those relating to the spatial appearance of objects, without consideration of their semantics, from the point of view of the agent. For example, one of three identical cubes may appear novel if it is displayed at a different angle or under different lighting conditions to the others. This could also be regarded as a subtype of contextual novelty.

Perceptual novelty refers to the way in which spatial elements of the visible scene may contrast with others due to their perceived differences, in particular due to their suddenness and saliency. Perceptual novelty therefore deals with comparisons of perceptual features rather than with comparisons of object, scene or event semantics: for example, a small object in a group of large objects, or a red object in a group of green objects would be highlighted for novelty due to saliency, while an object moving in the periphery could be deemed novel due to its eccentricity and velocity.

We incorporate a saliency map (Itti et al., 1998) for describing this type of spatial novelty (see (Chapter "Fundamentals of Agent Perception and Attention

Modelling”)), referring to it within the wider term of perceptual novelty. As has been pointed out in Itti and Baldi (2009), the saliency map can be thought of as novelty across space, while the novelty filters often described in robotics research could be considered as being novelty across time. In addition, this term also includes the concept of suddenness, which is related to perceptual influences such as the eccentricity at which a stimulus appears on the retina and its relative velocity. Perceptual novelty does not exclude contextual factors: for example, the saliency map considers spatial context in the sense of contrast of one spatial location with others, while visual similarity considers comparisons between all objects in the scene.

In addition to the general scene saliency, calculations regarding the visual similarity of objects may also be considered. A visual similarity module could be based on a perceptual grouping algorithm, such as that presented in Thorisson (1994). For a given scene rendering, this algorithm creates a graph where the nodes represent objects and the edges record the similarity between the objects with respect to their visual characteristics, such as colour, size and proximity. This process is based solely on visual input and does not utilise any semantic object information.

Suddenness may be related to the motion of objects, particularly in the periphery of vision. Eccentricity and relative velocity may also contribute to a notion of suddenness. The appearance of objects from nowhere, i.e. sudden onset, would also be one of a host of other factors which may contribute to the calculation of such a metric.

4.2.3 Partial Novelty

Partial novelty refers to when a known object type is encountered, but where that specific instance of the object has not been encountered before. This new instance of the object may contain sub-features that have not been observed previously in other instances of the same type. Partial novelty may be modelled by storing dynamic values for each object instance in the memory system.

An important concept here is habituation (Marsland et al., 2000), which relates to the reduction in behavioural response to a stimulus when it is perceived repeatedly. Unlike other forms of behavioural decrement, in the case of habituation, a change in the nature of the stimulus restores response back to its original level, a process known as dishabituation .

The memory model (Sect. 4.1.1) can be used to track objects in the environment that have been observed by the agent and the degree to which it has been familiarised with them through, e.g., repeated exposure and exploration. Objects falling within the gaze of the agent may be regarded as being attended to: as objects are attended to by the agent, the uncertainty level in memory corresponding to the object is reduced. As mentioned above, this value represents the completeness of the agent’s representation of the object in memory. Thus, as the agent attends to objects, their uncertainty value decreases. This value can be linked to habituation, so that those objects with low uncertainty values are regarded as having been habituated. Changes in an object’s state should lead to a corresponding increase in the uncertainty of the object, thus modelling a form of dishabituation.

4.2.4 Contextual Novelty

Contextual novelty refers to known objects occurring in unusual situations and implies a semantic comparison occurring between multiple objects. This type of novelty can also consist of matching the foreground elements in a scene with its predefined gist (Oliva and Torralba, 2001), such as an elephant in a street scene or a car in a fridge. Here, the ontology is necessary for calculating the semantic difference between objects. This operation could consist of a calculation for each object in the scene of what context it appears in now, and how many times this relationship has been previously encountered. An approach similar to that used in Oliva and Torralba (2001) may be adopted, where the scene is represented as a single entity with a vector of contextual features such as openness, naturalness. An important concept for applying this to virtual environments is that of scene partitioning, i.e. the splitting of the virtual environment into partitions. A scene ‘gist’ vector can be manually attributed to each zone in the environment. All objects lying within a certain zone can then be compared with the scene gist vector tagged for that zone to establish the novelty of the particular object in that scene context.

4.2.5 Semantic Novelty

Semantic novelty concerns the relationship between scene objects based on their categorisations according to the ontology. Approaches are available for computing semantics-based similarity decisions for ontologies (see, for example Chen and McLeod, 2005), so as to provide basic semantic relationships between objects present in the scene. The ontology is represented as a graph containing entities at nodes and their relationships as edges, as has been used by Navalpakkam and Itti (2002) for task relevance. Relationships may be of the following types: *is a*, *includes*, *part of*, *contains*, *similar* and *related*.

4.3 An Integrating Framework: Visual Attention

In addition to supporting components and a specification of the types of novelty, an important further consideration for constructing a system is how integration may occur between a novelty component and an agent framework. This section focuses on how all of the components can be integrated together, to form a coherent system capable of perception, decision and action. An integration concept described here is that of visual attention, where a number of models have been proposed or adapted for driving the attentive behaviours of computer agents based on visual external stimuli, to create what are referred to as ‘active vision’ (Aloimonos et al., 1987) for agents and robots: a more extended description can be found in (Chapter “Fundamentals of Agent Perception and Attention Modelling”). In this case, novelty detection may be one in which potentially relevant parts of the scene can be identified so that agents can adjust their attention appropriately towards new or otherwise important stimuli. The element we demonstrate here is part of the perceptual novelty category previously mentioned and consists of generating a saliency map signalling the conspicuity of constituent elements of the visual field (Itti et al., 1998).

4.3.1 Synthetic Perceptual Maps

The saliency map described in Sect. 4.2.2 is just one specific example of a number of different maps, based on the visual field of the agent, that may be used to represent the results of perceptual processing. A generalisation of these maps, referred to as synthetic perceptual maps (Peters, 2007), provides a way for different features to be represented, combined and operated on in a homogenous fashion based on the agent's synthetic perception (see Fig. 2). For the visual modality, synthetic perceptual maps are a virtual analogy of topographic retinotopic maps that represent the visual world as seen through the eyes of a viewer. They are rectangular, 2D gray scale maps corresponding to the agents field of view, where the value of a location in the map represents the strength of some particular feature or resultant operation based on the corresponding spatial location. Bottom-up saliency maps (Itti et al., 1998), task relevance maps (Navalpakkam and Itti, 2002) and the maps described here can all be viewed as instantiations of SPM's in our model. At each update of the agent's vision (see Sect. 4.1.2), two initial synthetic perceptual maps are created as inputs to the agent's perceptual pipeline. Two initial types of perceptual maps are created during each perceptual update: a full-scene map, based on a full colour rendering of the scene, and a false-colour map, based on a rendering of the scene where each object's colour is rendered uniquely without any extra operations, such as lighting or texturing, enabled. The saliency map can be viewed as a SPM created from the full-scene map. In the same way, this methodology allows for the modulation of the agents visual perception based on emotional (Adolphs, 2004), novel or even

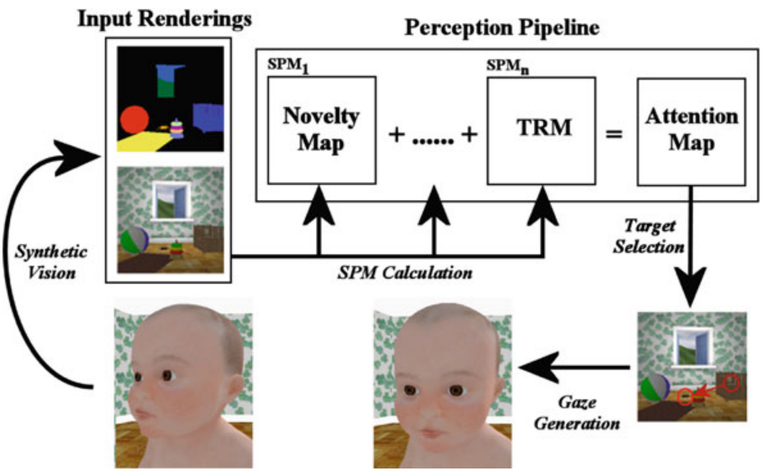

Fig. 2 Illustration of a perceptual pipeline for agents, creating a loop between perception and action, using novelty calculations to guide behaviour

memorised elements of the scene. For example, scene objects can be associated with threat values and based on an initial false-colour perceptual map, a specialised threat map can be created so that the strength in each element of the map is related to the threat value of the perceived object. This technique can also be used for establishing novelty based on objects' uncertainty values, creating a novelty master map.

4.3.2 Visual Perception Pipeline

The visual perception pipeline is a branched pipeline structure consisting of a number of different stages and partitions (see Fig. 2). Each stage in the pipeline contains a number of synthetic perceptual maps, each one the result of a processing operation on either the input maps or a map from a previous stage in the pipeline. SPM's can be combined together and the result of the perception pipeline is always a single master attention map, from which targets for the agent's gaze motions are selected.

The novelty master map can be further amalgamated with other features, the threat map previously mentioned, for example, in order to create a final master map, the attention map, which may be used for driving attentive behaviours. As with the other types of synthetic perceptual maps, the attention map is gray scale in nature; the intensities in the map represent the agent's attention to those respective parts of its visual field. An overall scene metric can also be evaluated, so that the novelty over the full scene can be used to drive explorative behaviours or have the agent engage in idle looking if the scene novelty falls below a certain threshold.

4.4 Outstanding Issues

The framework and components described here represent only one way in which a limited novelty detection system may be created for virtual agents. There are still many difficult theoretical and practical issues that remain to be addressed. A central issue involves the multitude of ways in which objects or events may be compared or considered to be novel from each other. No doubt there are many more super- and sub-categories to those that we have described here, and creating computational systems that can compare all dimensions of similarity in calculations seems an arduous, if not impossible task. Indeed, such categorisations must ultimately rest with the viewer, in terms of their perception of the scene, their prior experience, current mental and emotional state and task at hand. Assuming this is achieved, a further question then relates to how the weighting of these different dimensions should occur in order to construct a single overall novelty estimate, i.e. the novelty master map. Again, top-down biasing seems to be an important factor here, but how exact control should take place is still unclear. These issues make evaluation a difficult task. Additionally, the question of learning and adaptation of the system to new stimuli needs to be considered (Lungarella et al., 2003); the design presented here is far from flexible and is meant as one possible starting point for implementing a novelty capability. This capability can be framed as an early stage of an appraisal processing (Scherer, 2001) system for autonomous agents, by honing their attention to the relevant and important, supporting forgetting mechanisms and ultimately allowing them to better function in complex, unpredictable environments, real or virtual (Breazeal and Brooks, 2005).

Glossary

real novelty Refers to a new object instance that has never been or rarely been encountered before, although it may fall into a known object category. For example, a car may be deemed novel as this particualr instance has never been observed before.

perceptual novelty Perceptual novelty is the assembly of a new representation of an object never perceived in the past by the organism and requiring a new encoding in short-term and long-term memory.

partial novelty Partial novelty would be involved when an organism perceives an object looking like an already perceived object in the past but presenting some perceptual differences on one or several of these characteristics. For example if a baby used to play with a red ball, a new violet ball will not be totally new; its only one or several of these characteristic which are new; the categorization processes are very important in this concept of partial novelty.

contextual novelty Contextual novelty refers to the situations in which an object is perceived in a new context or emerged in a given stable context. For example a ball on the table, while such objects are usually on the floor, would induce a novelty detection phenomenon.

semantic novelty Semantic novelty refers to a situation in which the relationships between the objects or the concepts are organized in a new manner and have never been perceived such as in the past. The fact that individuals are able to create a new tool or a new concept from a series of well-known objects or ideas is characteristic of this type of novelty.

synthetic perceptual maps Gray scale topographic retinotopic maps that represent the visual world as seen through the eyes of a viewer, where the value of a location in the map is the strength of a feature or resultant operation based on the corresponding spatial location.

References

Adolphs R (2004, November) Emotional vision. Nat Neurosci 7(11):1167–1168

Aloimonos JY, Weiss I, Bandopadhay A (1987) Active vision in Proc. 1st Int. Conf. Comput. Vis., London, UK, pp 35–54

Ball SA (2004) Personality traits, disorders, and substance abuse. In: Stelmack RM (ed) On the psychobiology of personality: essays in honor of Marvin Zuckerman. Elsevier, New York, NY, pp 203–222

Bernstein AS (2002) The orienting response and stimulus significance: further comments. Biol Psychol 12(2–3):171–185

Breazeal C, Brooks R (2005) Robot emotion: a functional perspective. In: Fellous JM, Arbib MA (eds) Who needs emotions? The brain meets the Robot, Oxford University Press, NY, pp 271–310

Chen A, McLeod D (2005, May) Semantics-based similarity decisions for ontologies. In: ICEIS 2005, proceedings of the 7th international conference on enterprise information systems, Miami, USA, May

Chong H, Riis JL, McGinnis SM, Williams DM, Holcomb PJ, Daffner KR (2008) To ignore or explore: top-down modulation of novelty processing. J Cogn Neurosci 20(1):120–134

Cloninger CR, Svrakic DM, Przybeck TR (1993) A psychobiological model of temperament and character. Archiv General Psychiatry 50(12):975–990

Comerchero MD, Polich J (1999) P3a and P3b from typical auditory and visual stimuli. Clin Neurophysiol 110(1):24–30

Delplanque S, N'Diaye K, Scherer K, Grandjean D (2007) Spatial frequencies or emotional effects? A systematic measure of spatial frequencies for IAPS pictures by a discrete wavelet analysis. J Neurosci Methods 165(1):144–150

Desire L, Veissier I, Despres G, Boissy A (2004) On the Way to Assess Emotions in Animals: Do Lambs (Ovis aries) Evaluate an Event Through Its Suddenness, Novelty, or Unpredictability? J Comp Psychol 118(4):363–374

Friedman D, Cycowicz YM, Gaeta H (2001) The novelty P3: an event-related brain potential (ERP) sign of the brain's evaluation of novelty. Neurosci Biobehav Rev 25(4):355–373

Friedman D, Simpson G, Hamberger M (1993) Age-related changes in scalp topography to novel and target stimuli. Psychophysiology 30(4):383–396

Grandjean D (2005) Etude lectrophysiologique des processus cognitifs dans la gense de l'motion. University of Geneva, Geneva

Grandjean D, Sander D, Scherer KR (2008) Conscious emotional experience emerges as a function of multilevel, appraisaldriven response synchronization. Conscious Cogn 17(2):484–495

Grandjean D, Scherer K (2008) Unpacking the cognitive architecture of emotion processes. Emotion 8(3):341–351

Habib R, McIntosh AR, Wheeler MA, Tulving E (2003) Memory encoding and hippocampally-based novelty/familiarity discrimination networks. Neuropsychologia 41(3):271–279

Halgren E, Baudena P, Clarke JM, Heit G, Liegeois C, Chauvel P, Musolino A (1995a) Intracerebral potentials to rare target and distractor auditory and visual stimuli. I. Superior temporal plane and parietal lobe. Electroencephalogr Clin Neurophysiol 94(3):191–220

Halgren E, Baudena P, Clarke JM, Heit G, Marinkovic K, Devaux B, Vignal JP, Biraben A (1995b) Intracerebral potentials to rare target and distractor auditory and visual stimuli. II. Medial, lateral and posterior temporal lobe. Electroencephalogr Clin Neurophysiol 94(4): 229–250

Henson RN, Rugg MD (2003) Neural response suppression, haemodynamic repetition effects, and behavioural priming. Neuropsychologia 41(3):263–270

Hodge VJ, Austin J (2004) A survey of outlier detection methodologies. Artif Intell Rev, 22(2): 85–126

Itti L, Baldi P (2009) Bayesian surprise attracts human attention. Vis Res 49(10):1295–1306

Itti L, Koch C, Niebur E (1998 November) A model of saliency-based visual attention for rapid scene analysis. IEEE Trans Pattern Anal Mach Intell (PAMI) 20(11):1254–1259

Knight RT (1996) Contribution of human hippocampal region to novelty detection. Nature, 383(6597):256–259

Knight RT (1984) Decreased response to novel stimuli after prefrontal lesions in man. Electroencephalogr Clin Neurophysiol 59(1):9–20 (1984)

Knight RT, Nakada T (1998) Cortico-limbic circuits and novelty: a review of EEG and blood flow data. Rev Neurosci 9(1):57–70 (1998)

Lang PJ, Bradley MM, Cuthbert BN (2005) International affective picture system (IAPS): affective ratings of pictures and instruction manual. Technical Report A-6. University of Floridao, Gainesville, FL

Lazarus R (1991) Emotion and adaptation. Oxford University Press New York, NY

Leventhal H, Scherer K (1987) The relationship of emotion to cognition: a functional approach to a semantic controversy. Cogn Emot 1(1):3–28

Lungarella M, Metta G, Pfeifer R, Sandini G (2003) Developmental robotics: a survey. Connection Sci 15(4):151–190

Markou M, Singh S (2003) Novelty detection: a review, part 2: neural network based approaches. Signal Process 83(12):2499–2521

Marsland S, Nehmzow U, Shapiro J (2002) Environment-specific novelty detection. In: From animals to animats, Proceedings of 7th international conference on simulation of adaptive behaviour, Edinburgh

Marsland S, Nehmzow U, Shapiro J (2000 January) Detecting novel features of an environment using habituation. In: Proceedings of simulation of adaptive behaviour. MIT Press, Cambridge, MA, pp 189–198

Naatanen R, Paavilainen P, Rinne T, Alho K (2007) The mismatch negativity (MMN) in basic research of central auditory processing: a review. Clin Neurophysiol 118(12):2544–2590

Navalpakkam V, Itti L (2002, November) A goal oriented attention guidance model. Lect Notes Comput Sci 2525:453–461

Neto HV, Nehmzow U (2007) Visual novelty detection with automatic scale selection. Robot Auton Syst 55(9):711–719

Neto HV (2006) Visual novelty detection for autonomous inspection robots. PhD thesis, University of Essex, Colchester, UK

Oliva A, Torralba A (2001, May) Modeling the shape of the scene: A holistic representation of the spatial envelope. Int J Comput Vis 42(3):145–175, May

Peters C (2007) Designing an emotional and attentive virtual infant. In: Proceedings of the 2nd international conference on affective computing and intelligent interaction (ACII), Lisbon, Portugal, September 12–14, 2007. Lect Notes Comput Sci (LNCS) 4738:386–397

Peters C (2006, September) Designing synthetic memory systems for supporting autonomous embodied agent behaviour. In: Proceedings of the 15th international symposium on robot and human interactive communication (RO-MAN), University of Hertfordshire, Hatfield, UK, pp 14–19

Peters C, Grandjean D (2008, May) A visual novelty detection component for virtual agents. In: Paletta L (ed) Proceedings of the fifth international workshop on attention and performance in computational vision (WAPCV), Santorini, Greece, pp 289–300

Polich J (1988) Bifurcated P300 peaks: P3a and P3b revisited? J Clin Neurophysiol 5(3):287–294

Polich J (2007) Updating P300: an integrative theory of P3a and P3b. Clin Neurophysiol 118(10):2128–2148

Ranganath C, Rainer G (2003) Neural mechanisms for detecting and remembering novel events. Nat Rev Neurosci 4(3):193–202

Ringo JL (1996) Stimulus specific adaptation in inferior temporal and medial temporal cortex of the monkey. Behav Brain Res 76(1–2):191–197

Sander D, Grandjean D, Scherer KR (2005) A systems approach to appraisal mechanisms in emotion. Neural Netw 18(4):317–352

Savage-Rumbaugh ES, Toth N, Schick K (2003) Kanzi Learns to Knap Stone Tools. In: Washburn DA (ed) Primate perspectives on behavior and cognition. American Psychological Association Washington, DC, pp 279–291

Scherer K (2001) Appraisal processes in emotion: theory, methods, research, chapter Appraisal considered as a process of multilevel sequential checking. Oxford University Press, New York, NY, pp 92–120

Scherer KR (1984) On the nature and function of emotion. A component process approach. In: Scherer KR, Ekman P (eds) Approaches to emotion. Erlbaum, Hillsdale, pp 293–317

Siddle DAT, Lipp OV (1997) Orienting, habituation, and information processing: The effects of omission, the role of expectancy, and the problem of dishabituation. In: Lang PJ, Simons RF, Balaban MT (eds) Attention and orienting: sensory and motivational processes. Lawrence Erlbaum Associates, Mahwah, NJ, pp 23–40

Thorisson K (1994) Simulated perceptual grouping: an application to human computer interaction. In: Proceedings of the 16th annual conference of cognitive science society, Atlanta GA, pp 876–881

Tulving E, Kroll N (1995) Novelty assessment in the brain and long-term memory encoding. Psychonomic Bull Rev 2(3):387–390

Tulving E, Markowitsch HJ, Kapur S, Habib R, Houle S (1994) Novelty encoding networks in the human brain: positron emission tomography data. Neuroreport 5(18):2525–2528