

Marc Cavazza*
Jean-Luc Lugin
School of Computing
University of Teesside
Middlesbrough, TS1 3BA, UK

Marc Buehner
School of Psychology
Cardiff University
Cardiff, DF10 3AT, UK

Causal Perception in Virtual Reality and its Implications for Presence Factors

Abstract

Causality is an important aspect of how we construct reality. Yet, while many psychological phenomena have been studied in their relation to virtual reality (VR), very little work has been dedicated specifically to causal perception, despite its potential relevance for user interaction and presence. In this paper, we describe the development of a virtual environment supporting experiments with causal perception. The system, inspired from psychological data, operates by intercepting events in the virtual world, so as to create artificial co-occurrences between events and their subsequent effects. After recognizing high-level events and formalizing them with a symbolic representation inspired from robotics planning, it modifies the events' effects using knowledge-based operators. The re-activation of the modified events creates co-occurrences inducing causal impressions in the user. We conducted experiments with fifty-three subjects who had to interact with virtual world objects and were presented with alternative consequences for their actions, generated by the system using various levels of plausibility. At the same time, these subjects had to answer ten items from the Presence Questionnaire corresponding mainly to control and realism factors: causal perception appears to have a positive impact on these items. The implications of this work are twofold: first, causal perception can provide an interesting experimental setting for some presence determinants, and second, the elicitation of causal impressions can become part of VR technologies to provide new forms of VR experiences.

I Introduction

Causality is an essential aspect of our common sense understanding of the physical world, whose implications for our sense of reality have been discussed by physicists (d'Espagnat, 2002) and philosophers (Zubiri, 2001) alike. In cognitive psychology, causal perception has been considered as one of the main phenomena through which we perceive our everyday reality. This naturally makes causality an interesting focus of experimentation for virtual reality.

Although rarely referred to explicitly, there is significant evidence of the use of causality in presence research, most specifically when considering those aspects of presence dealing with action, agency, environment control, and the realism of the environment's responses. From a fundamental perspective, this

should not be entirely surprising, as causality is one of the few psychological phenomena bridging the gap between perception and high-level cognitive concepts (Scholl & Tremoulet, 2000). One of the early works that introduced concepts related to causal perception was that of Loomis (1992) on distal attribution, although causality was not considered explicitly. Mantovani and Riva (1999) following Schloerb (1995) introduced the concept of *causal interaction* as an essential aspect of presence. In Witmer and Singer's Presence Questionnaire (Witmer & Singer, 1998), several items typically involve causal perception. For instance, Item #2 of their original questionnaire reads "How responsive was the environment to actions that you initiated?" Further, their use of McGreevy's argument (McGreevy, 1992) about "continuities, connectiveness and coherence of the stimulus flow" is also evocative of causal perception. Finally, Zahorik and Jenison (1998) in their in-depth discussion of the phenomenological conditions of presence, advocate that a "lawful response" from the environment to our actions should be a major determinant of presence. More importantly, their strong Gibsonian perspective is somehow close, in terms of its philosophy of perception, to that of Michotte, the father of causal perception (Michotte, 1963).

In this paper, we advocate that causal perception can play a major role in the subjects' experience in VR. We describe experiments in causal perception that investigate the connections between causality and several determinants of presence, in particular those identified in terms of interactivity and environment controllability.

While there exists a rich literature on various aspects of perception in VR, there has been little work specifically investigating causal perception in VR. The only specific studies of techniques for enforcing causality have taken place in distributed virtual environments (see, e.g., Roberts & Sharkey, 1997), and has investigated the correct propagation of consequences rather than the fundamental determinants of causality itself, which we are discussing here.

The only specific work on the visualization of causal relations is that of Ware et al. (1999), which, however, addresses 2D noninteractive visualization, and as such

does not directly transpose to the study of presence in virtual environments. Finally, in the field of experimental psychology, the use of 3D animation has previously been reported to elicit causal perception in subjects (Wolff & Zettergren, 2002). However, these systems were noninteractive 3D animations: only virtual environments can support proper experiments on presence factors such as control.

This is why we have developed specific techniques to control causal impressions in virtual reality, which we have used to support our experiments.

In the next sections, after introducing the determinants of causal perception, we describe our approach to the elicitation of causal perception, which is based on the explicit representation of actions taking place in a virtual environment and their consequences. We devise a test environment supporting experiments with subjects that simultaneously investigate causal perception and certain determinants of presence. We analyze these findings and suggest new directions for the use of causality in VR, both as a theoretical framework and as a support for the development of VR technology.

2 Causal Perception

In the context of this research, we should only consider causality between physical events, sometimes referred to as *event causality* (Galavotti, 2001), to distinguish it from other causal relations established between properties, such as "smoking causes diseases" (see Pearl, 2000; or Cheng, 1997; for authoritative reports on property causality).

This distinction can also be rephrased as one between causal induction and causal perception. The former is based on a Humean perspective of causality, according to which causal relations cannot be perceived directly. Instead, causal knowledge has to be derived from observable entities such as the presence and absence of events, and the spatio-temporal relations between them. One of Hume's premises for causal induction is the "constant conjunction" of the events in question; in other words, only if one is exposed to multiple cause-effect sequences will one form a mental causal link be-

tween them. There are various psychological theories about how such causal *learning* takes place (see Cheng, 1997, or Buehner & Cheng, 2005, for an overview).

Causal perception is in stark contrast to induction, in that it appears to be spontaneous, irresistible, and not dependent on multiple trials. Michotte (1963) was the first to note that the apprehension of causality often appears perceptual in nature, despite causality being generally considered a high-level property of the world (Scholl & Nakayama, 2002). In Michotte's theory, causality is attributed to co-occurring events from their spatio-temporal contiguity. The canonical example used by Michotte consists of one pool ball striking another, thus "launching" the latter, which acquires autonomous motion. In these experiments, Michotte explored the spatio-temporal conditions (namely, the gap between the two balls and the delay between the two motion events) for the attribution of causality between the two events. Central to the attribution of a causal relation between the two events is the time interval separating them: Michotte (1963) reported that if the second event was delayed by more than 150 ms, 50% of subjects failed to recognize a causal relation between them. Recently, White and Milne (1999) have extended Michotte's catalog of functional relations, which suggests that causal perception is not restricted to launching events.

The distinction between perceptual causality and causal judgment or induction has been investigated by Schlottmann and Shanks (1992), who found an interesting dissociation: Judgment and perception can reach opposite conclusions, with participants *expecting* an effect (thus knowing that the cause will produce it), yet reporting that the interaction looked non-causal due to a spatial or temporal gap. This need not mean that causal induction and causal perception are two entirely independent processes. Cheng (1993), for instance, provided an inductive analysis of Michotteian perceptual causality. Recent work by Scholl and Nakayama (2002) likewise suggests that the perceptual system, despite its immediate recognition of causal stimuli, incorporates an inductive component.

In sum, causal perception lies at an interesting intersection between perception and cognition, as if a high-level property of visual events can sometimes be per-

ceived directly (Scholl & Nakayama, 2002). This reinforces its interest in the current context, as one of the major difficulties of presence studies has precisely been to relate perceptive aspects to higher-level psychological determinants (see, e.g., Riva, Davide, & Ijsselstein, 2003).

Our objective is to induce causal perception from collision events in a virtual world. We hypothesize that causal perception can be elicited through the real-time manipulation of events as they take place, to create new co-occurrences that can be interpreted in causal terms by a user. Figure 1 illustrates this process: the fall of the glass on the table is immediately followed by the nearby menu tilting.

Because of the importance of temporal aspects for causal perception, this whole process of creating co-occurrences should be inserted in the natural interaction cycle of the VR system (including in terms of sampling rate), so as to be transparent to the user.

3 System Architecture and Overview

We have developed a "causal engine," which operates by modifying virtual world events as they take place, so as to create event co-occurrences of the kind pictured on Figure 1 that will induce causal impressions in the user. The causal engine operates by intercepting ongoing events and altering them while their effects are temporarily "frozen."

What the causal engine exploits is the strong propensity of humans to perceive co-occurring events as causally linked, especially when they initiate the first event through their own actions. From the subject's perspective, his or her interactions with the world objects will not result in their ordinary consequences. Rather, these default consequences will be "intercepted" and substituted with other effects. For instance, while a glass falling on a table would normally shatter (spilling its contents), the causal engine can generate alternative effects, such as the glass landing intact on the table but causing another glass to tumble and spill its contents.

Our causal module has been developed as an additional

layer on top of the visualization engine, the Unreal Tournament 2003 Game Engine™ (Lewis & Jacobson, 2002), corresponding to over 70,000 lines of C++ and UnrealScript code. The causal layer communicates with the visualization engine using UDP messages and is integrated to the engine's behavioral API using a set of dedicated UnrealScript wrappers. The software architecture (Figure 1) was designed to support the system's main operating cycle, which consists of i) intercepting native events at a regular sampling rate, ii) recognizing the high-level actions they correspond to, and iii) modifying their default consequences, so as to create new co-occurrences on a principled basis.

As such, it integrates the following components: an Event Interception System (EIS), an action recognition layer, and a causal layer.

3.1 Event Interception System

An Event Interception System (EIS) was developed on top of the native event system in UT 2003. Event systems are generally derived from the low-level graphical primitives for collision detection (between objects, and between objects and volumes). Events constitute an intermediate-level formalism for interaction and play an important role in the basic architecture of VR systems (Jiang, Kessler, & Nonnemaker, 2002). This module processes occurrences of the game engine's native events, to produce a refined set of low-level events henceforth called Basic Events (BE). For instance, the magnitude of the object momentum in a colliding event can be used to instantiate a `Hit(?obj, ?surface)` BE, from system-level events signaling the collision. This makes it possible to differentiate between various types of impact (hitting, landing, etc.).

The EIS derives its name from the fact that it overrides part of the native physics engine, namely the processing of physical interactions between objects, corresponding to various kinds of contact and collisions.

3.2 Action Recognition Layer

An action recognition layer interprets BE in terms of higher-level actions. Our representation for high-level

actions is referred to as CE (for "Cause-Effect") structures, and constitutes a detailed formalization of actions and their consequences. For instance, the `Break-on-impact` CE represented in Figure 2 describes the action through which a fragile object can be broken when hitting (or being hit by) a harder one. The CE representation contains in its `condition` part the physical properties of its objects, and describes the action's consequences in its `effect` part. This representation is thus organized around the notion of action and consequences, which is the most appropriate to represent event causality. A detailed description of the formalism is given in Section 4.1.

This module constitutes one of the innovative aspects of our approach, as it underpins the virtual world with an ontology of possible actions that could involve/affect world objects. This is somehow similar to the conceptions put forward by Mantovani and Riva (1999).

During a sampling cycle (5–10 ms), it operates by receiving a stream of BE intercepted during the current time sample, and outputs a set of instantiated CE to the causal engine, which correspond to the recognized actions. It should be noted that these CE are "frozen" prior to their processing by the causal engine, that is, their `effect` part is prevented from being executed.

3.3 Causal Engine

The "causal engine" itself produces new event co-occurrences in the virtual world by modifying the effects of the instantiated CE to generate alternative effects. These event co-occurrences will, in turn, induce causal impressions of the kind discussed above. The modification of the CE's effects is the key mechanisms for generating such co-occurrences.

This module receives a set of CE from the action recognition system and (after altering their default effects) outputs a set of modified CEs, which are ready to be reactivated, and have their effects triggered in the virtual world.

The detailed behavior of these modules, as well as the working cycle of the whole system, is discussed in subsequent sections.

4 The Causal Engine

The causal engine is developed on top of the native event system, which supports event discretization in the physics engine (which in the case of UT 2003 is the Karma™ engine). This means that the default physics is overridden by the new mechanisms provided by the causal engine for any event involving interactions between objects (other aspects such as kinematics are not altered and remain under the control of the physics engine).

The rationale for using an explicit representation for actions is that it supports action modification on a principled basis. Rather than directly associating effects to specific actions, which can be a tedious process and requires a priori definition of that association, the action representation can be modified using high-level operations, for instance, substituting the action's object or modifying its effects. Only this can support proper experimentations with causality, as it makes it possible to explore alternatives in a systematic fashion, and suggests associations to be experimented with.

The causal engine is a generic system that operates regardless of the action's origin (i.e., a user intervention, a simulated process, etc.). However, for the purpose of these experiments we have only considered actions initiated by the user. We are also only concerned with effects that could be perceived as directly resulting from the user's actions rather than complex causal chains. Even in the case where effects tend to propagate further (as in the case of an impact causing a glass to tilt, and consequently to spill out its contents), only the first consequence (the glass tilting) will be subject to modification by the causal engine. Further consequences within the same time sample are not altered.

The causal engine operates continuously through sampling cycles that are initiated by the occurrence of actions in the virtual world. Basically, the occurrence of events affecting world objects initiates a sampling cycle, during which the system recognizes potential CE and stores them while inhibiting their effects (it could be said that it freezes them). The causal engine then transforms these frozen CE, by altering their effects, before reactivating them. This reactivation then initiates a new sampling cycle.

4.1 Action Representations and the Creation of Co-Occurrences

Action representation should be supported by an appropriate formalism for change-inducing events, which should clearly identify actions as well as their consequences. The second step consists in defining an ontology of such events, that is, describing the most important high-level events that can be recognized in the virtual environment. This form of representation echoes the proposal of Mantovani and Riva (1999), according to which a virtual environment should be characterized by an ontology, where, however, we consider an ontology for actions.

Typically, an action is represented in the causal engine using our Cause-Effect action formalism (henceforth CE) inspired from those used in planning and robotics, such as STRIPS (Fikes & Nilsson, 1971). These representations originally describe operators transforming states of affairs in the world. They tend to be organized around pre-conditions, that is, conditions that should be satisfied for that action to take place, and post-conditions, that is, those world changes induced by their application.

Our CE formalism is composed of three main fields, which are analogous to the SIPE representation (Wilkins, 1998); we shall illustrate it in an example that will play an important role in our experiments, the *Break-on-impact* CE, which describes the event by which a fragile object will shatter upon colliding with a hard object (Figure 2).

The first field, called *trigger*, contains the basic events from which the CE can be recognized, and which prompts instantiation of the corresponding CE representation. In the *Break-on-impact* CE, this field contains the BE *Hit*, derived from the low-level event systems of the UT 2003 engine, such as *Touch*, *Bump*, *Kimpact* events. Any occurrence of a *Hit* BE can potentially activate the instantiation of a *Break-on-impact* CE (Figure 2, event 1).

The *condition* field determines the physical properties that should be satisfied by objects taking part in that specific action, such as being “movable” or “breakable” (these semantic properties being characterized by

physical properties). For instance (Figure 2), a moving object hitting another object will break only if its substance is fragile and the object it hits is hard (Figure 2, event 2).

The `condition` field is used to filter between candidate CEs primed by a similar basic events. For instance, the CE representations for `Bouncing` and `Break-on-impact` can be activated from the same `Hit` BE. It is the physical properties of the objects involved in the action that will determine which candidate CE describes the situation at hand.

The `trigger` and `condition` fields govern the CE instantiation: once these fields can be instantiated by the CE recognition mechanism, a corresponding CE representation is created, whose modification will create event co-occurrences. It can be noted that CE actually represent events rather than intentional actions, as they ignore the course of events preceding the CE. The above CE for `breaking-on-impact` is recognized in a similar fashion regardless of the origin of the impact, whether the object is falling on a hard surface, has been launched, or is struck by a harder object. This contributes to making this kind of representation more generic and expressive.

Finally, the `effect` field corresponds to the consequence part of the CE. It contains the default transformations to be applied to the objects affected by the CE (Figure 2, event 3). For instance, consider a case of a glass shattering, the deletion of the glass and the creation of glass fragments (Figure 2, event 4). The CE formalism plays a central role in the creation of event co-occurrences. As it essentially associates actions with their consequences, it can be modified, for instance, by substituting alternative consequences to the default ones. In that sense, new event co-occurrences are produced by i) inhibiting the activation of CE's default effects immediately after their instantiation, ii) modifying the CE's effects while these are suspended, and iii) re-activating the CE's modified effects. As a result, for any given CE processed by the system, the user would perceive the corresponding triggering event followed by an alternative effect (see Figure 3 for a set of alternative effects following the fall of an empty glass on a table).

4.2 The Causal Engine in Action

Throughout this discussion, we will illustrate the behavior of the causal engine in a simple example (Figure 3), to be part of our subsequent experiments. The test case we will consider is that of a glass being grasped, then dropped by the user from a given height onto the surface of a table, which can also hold other objects, such as similar glasses, a cardboard menu, a candle, and a beer bottle. The default physical behavior would consist of the glass shattering on impact (Figure 2), as it would be directly obtained from the use of the system's native physics engine. This type of event also constitutes a traditional example for causality (see, e.g., Collins, Hall, & Paul, 2004). The corresponding CE, `Break-on-impact`, is represented in Figure 2: its triggering event is the glass hitting the table surface.

The causal engine operates by first instantiating CE representations for ongoing actions, then modifying the effects of some or all of these representations. CE instantiation is a bottom-up process, which starts with the processing of incoming Basic Events. BE are directly generated by virtual world objects taking part in interactions. The native event system has been modified to recognize a specific set of fifteen Basic Events (such as `Hit`, `Push`, `Aligned`, `Fall`, etc.). Our event controller continuously intercepts BE instances as they are generated, as shown in Figure 2: when the glass hits the table, this generates a `Hit (glass # 1, table # 1)` BE.

The next step consists of the recognition of potential actions through the identification of CE, which is a discrete process executed at a high sampling frequency, typically 5 to 10 ms. During this operation, the intercepted BE instances are redistributed to relevant CE prototype representations (Figure 2, event 1), that is, those CE that have a compatible BE in their `trigger` field. In our example, the `Hit (glass # 1 table # 1)` BE activates several CE, among which are the `Break-on-impact (glass # 1 table # 1)`. This step is optimized through off-line pre-processing, which indexes CE on their triggering BE categories. Once activated by the reception of a BE, the CE become candidates for instantiation.




Figure 1. System overview.

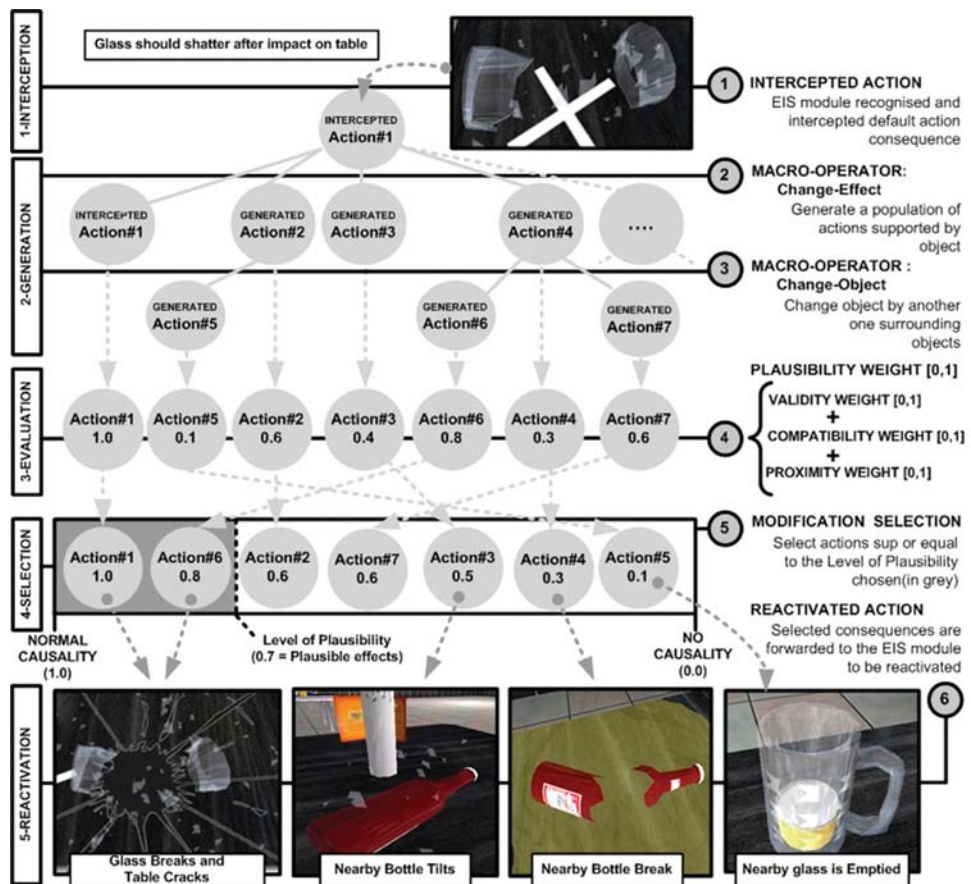


Figure 3. Level of plausibility and the action generation algorithm.




Figure 2. The cause-and-effect (CE) formalism and its use of objects' semantic properties.

Activated CE representations immediately execute their Condition predicates on those BE that have activated them. If objects involved in a BE instance satisfy the entire set of predicates, a CE instance is generated. In the falling glass example (Figure 2, event 2), Substance (glass # 1, FRAGILE) is true, and so is Substance (table # 1, HARD). Hence a Break-on-impact (glass # 1, table # 1) CE is instantiated. Once a CE is generated, it becomes a target for transformation by the causal engine. The CE recognition system operates essentially within the visualization engine, while the causal engine is an external module communicating via UDP. The CE instances are thus transmitted to the causal engine encoded as a set of data frames. Once it receives a set of instantiated CE, the causal engine modifies the action's outcome by altering the effect field of the CE representation (Figure 2, event 3). The effect field contains an action (e.g., shattering) to be applied to the CE's objects (e.g., the falling glass). This is why the causal engine can modify either the type of action, or the objects affected, or both. It does so by applying specific trans-

formation operators (called Macro-Operators or MOP). For instance, the change_object MOP replaces the CE's object with another virtual world object. Its consequences are visible in Figure 3. The default object of the Break-on-impact CE is the falling glass, which should shatter on landing. The causal engine intercepts that event and substitutes the Break-on-impact object with another one, which is table surface. As a result, the falling pint glass lands on the table, "causing" the other pint glass to shatter. From the user's perspective, the normal cause-effect sequence is disrupted: the triggering event of a given CE, in this case the glass falling on a table, will be followed, not by its default consequence (e.g., the falling glass breaking), but by an alternative effect (e.g., a nearby glass breaking without being directly hit). From an identical initial context (Figure 5, event 1), the same Change_object MOP could have associated the table rather than the glass to the CE effect (Figure 5, event 3). This results in the table, rather than the glass, breaking upon impact, even though it is by default the hardest object (Figure 5, event 4). This also illustrates the use of generic proce-

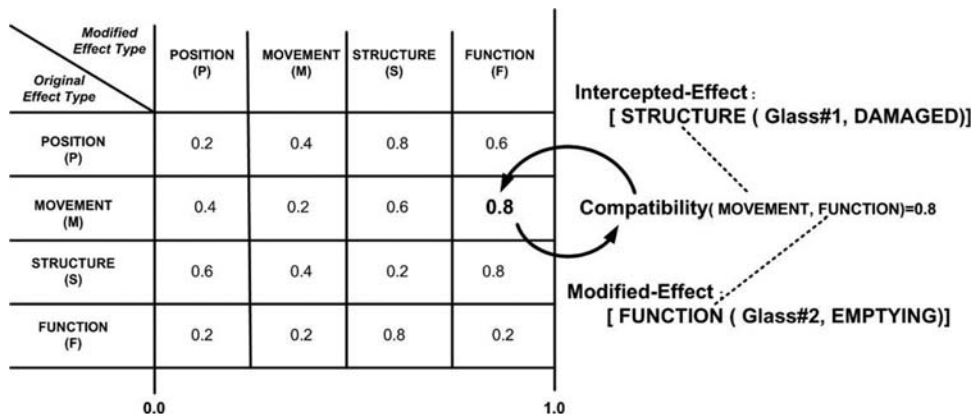


Figure 4. The compatibility matrix provides heuristics based on the difference in type between the original effect and the modified one.

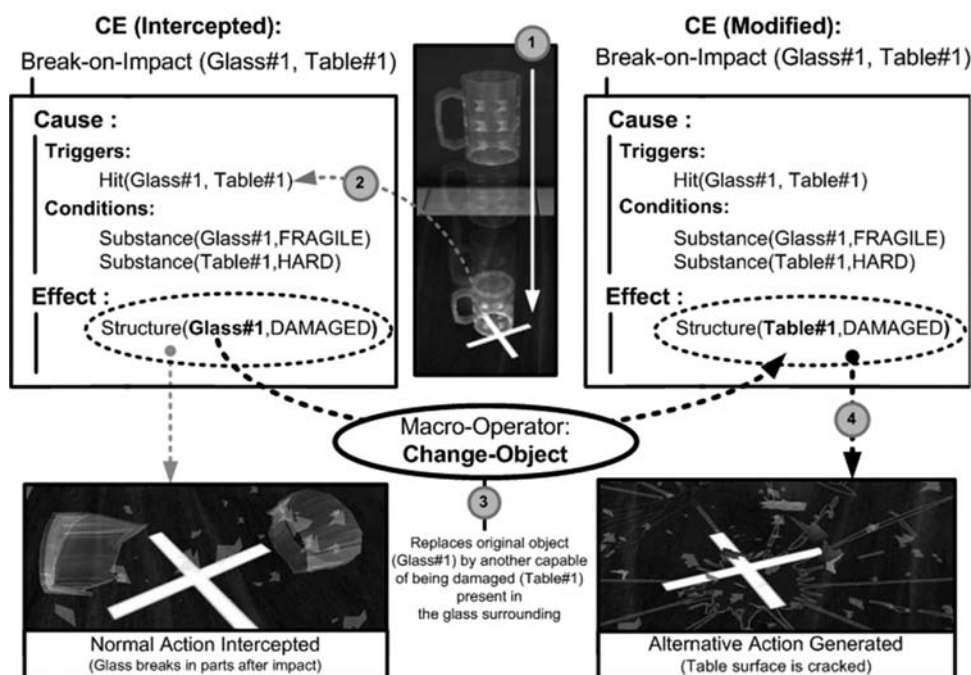


Figure 5. Example of macro-operator application (creation of an artificial co-occurrence).

dures for effects, which depend on generic physical properties, associated with specific animation visualizing object transformations. For instance, the generic state DAMAGED is automatically translated into *shattering* when applied to a glass, while, when applied to the table on which an object falls, it corresponds to the *cracking*

of the table surface (Figure 5, event 4). As explained in the following section, transformations involving object substitutions are based on semantic measures of action and object compatibility. These measures can be tuned to favor substitution of objects having different physical properties, so as to induce different levels of believability.

4.3 Alternative Action Generation

From the modifications of an intercepted action, the causal engine generates alternative consequences and evaluates their plausibility with regard to the normal effects expected. As shown in Figure 3, simple variations of a single parameter, called “level of plausibility,” supports the generation of different consequences that vary from normal to plausible, implausible, or completely unrealistic effects. The mechanism allowing such flexibility relies on a heuristic search executed by the causal engine, when an action in progress is intercepted by our EIS module (Figure 3, event 1). The search is based on three successive processes: *Generation*, *Evaluation*, and *Selection*.

The *Generation* phase (Figure 3, event 2) creates a collection of potential consequences based on the population of objects surrounding the initial action location. The algorithm successively applies a list of Macro-Operators on the intercepted action. For our experiment, the algorithm is using two different MOPs: *Change-Effect* and *Change-Object*. As previously shown (Figure 5), those transformation operators create alternative actions by directly modifying the *EFFECT* type or the object of the action considered. Here, *Change-Object* replaces the *Glass* object instance (*Glass # 1*) by the *Table* instance (*Table # 1*), its closest object. Hence, if the modified action is reactivated, the table will suffer damage instead of the glass. In a similar fashion, the *Change-Effect* operator modifies the effect type of an action, replacing the default effect by another one supported by the object (Figure 3, event 3). For instance, the effect by which the structure of the glass is damaged (*Structure (Glass # 1, DAMAGED)*) could be replaced by a simple change in the glass position (*Position (Glass # 1, TILTED)*), or *Movement (Glass # 1, REBOUNDS)*). The successive combination of those two operators quickly develops a large set of possible effects around the initial action.

The *Evaluation* phase is next. Once the list of applicable MOPs has been executed, the collection of action generated along the previous phase is evaluated in terms of plausibility. As shown in Figure 3, event 4, our sys-

tem associates to each action a normalized Plausibility weight, which is calculated from three heuristic values named validity, compatibility, and proximity weight.

$$\text{Plausibility Weight} = \text{Validity Weight} \\ + \text{Compatibility Weight} + \text{Proximity Weight}$$

The *validity weight* simply considers if a generated action is actually applicable with an object by testing if the object, previously substituted by the MOP, satisfies the action’s conditions. A validity weight of 1.0 means the entire set of pre-conditions is satisfied. By contrast, a value of 0 means that no object properties can satisfy the preconditions.

The *compatibility weight* is computed using a matrix associating a heuristic value to each possible combination of effect type (Figure 4). For instance, changing a *MOVEMENT* effect, such as *Tilting*, by another *MOVEMENT* effect, such as *sliding*, appears a lot less disruptive than replacing *Tilting* by a *FUNCTION*-type effect such as *Emptying*. The Plausibility matrix has been initially established by identifying analogies between potential consequences of a sample set of actions. In a subsequent step, the weights associated with the matrix elements have been readjusted according to feedback from previous user experiments.

However, the plausibility of a modification is also considerably influenced by the spatial contiguity of co-occurring events. Based on the principle that correlation between distant events is less likely to induce causal perception, we have reinforced the plausibility weight by spatial consideration using a proximity weight. A spatial weight valued at 0.0 corresponds to the original object’s position, where a value of 1.0 represents the farthest object.

In the *Selection* process (Figure 3, event 5) our level of plausibility is acting as a threshold to guide the search process toward a plausible or implausible consequence according to its value. A level of plausibility of 1.0 corresponds to the default consequence, while a zero value represents a total absence of effect. The range of alternative action is then classified in decreasing order of Plausibility weight. As a first step, the process extracts

the set of actions with a value equal or superior to the level chosen. Then, for each object involved in this pre-set, it opts for action closer to the plausibility level wished. Finally, the set of selected action will be forwarded to the EIS module to be reactivated (Figure 3, event 6).

The combination of MOp and heuristic search provides an original approach that allows a systematic exploration of a transformation space. The innovative aspect of our system lies in the possibility of adjusting dynamically the plausibility of an action's output using a single variable without specifying manually the consequences of a given action in a given context (which would be more likely to introduce a bias).

Overall, the system performance is in line with its initial design constraints, which imposed a response time in the order of 90–120 ms. The action recognition and reactivation process is achieved between 40–60 ms, while the action modification occurs in a range of time of 20–60 ms.

4.4 Results

Figure 3 shows some co-occurrences generated by the causal engine used in our experiments, for different tunings of the “realness” heuristic. In all realistic configurations, the falling glass would shatter on impact, but the system would generate additional effects, resulting in some of the example associations depicted in Figure 6 (for instance, the falling glass will shatter upon impacting on the table, with additional effects affecting one neighboring object; e.g., the cardboard menu falls). It should thus be noted that realistic effects do not result from the use of the default's physics engine: they are still artificial effects created by the system which are simply plausible, plausibility being defined as the preservation of basic mechanisms (the falling glass still shatters) and physical compatibility of effects (e.g., motion effect like tilting is more similar to a motion effect like sliding). The definition of implausible causality rests on the creation of event co-occurrences for which no straightforward relations exist between events. For instance, upon impact of the falling glass on the table, another glass will start emptying itself without its walls being cracked or the glass tilting at all. A fundamental aspect

of the system is that these co-occurrences can be automatically generated by tuning one single parameter in the heuristic search: the level of plausibility targeted. The actual generation of co-occurrences is thus dynamic and independent from the experiment's design, leading to the production of variable effects (both in nature and order) for each user.

A preliminary evaluation of the causal engine should compare its performance (in terms of response times) with data from psychological literature prior to the design of causal perception experiments involving subjects. Our tests have shown the overall response time to be on average 90–100 ms. This has to be compared to data in the psychological literature, which reported the maximum delay between consecutive events for these to be perceived as causally linked. In the original experiments from Michotte (1963), events delayed by more than 150 ms progressively ceased to be perceived as causally linked. Buehner and May (2003) contrasted “immediate” and “delayed” action-outcome sequences. The average response time on immediate pairings was “less than 0.25 s” (Buehner & May, 2003, p. 884) and participants assessed action-outcome contingencies under such a schedule accurately. Finally, when interpreting Michotte-style launching events, Kruschke and Fragassi (1996) considered that motion ampliation (considered to account for causal impressions in Michotte's theory) took place within a critical 200 ms interval. All these data suggest that the system's response time is compatible with results from the psychological literature: as a consequence, the co-occurrences generated should be perceived by the vast majority of subjects as sufficiently close to induce causal perception.

5 The Experiments

Our experiments aimed at evaluating the possible association between causal perception and some previously described factors of presence (Witmer & Singer, 1998), mostly described as *control factors*. In that sense, the variable we controlled was the elicitation of causal perception, while measuring presence factors by using a 10-item subset of the Presence Questionnaire. We have

selected 10 questions from the original Presence Questionnaire of Witmer and Singer, considering them from the perspective of how they could relate to simulated causality. This questionnaire is made up of 9 of the 12 questions of the Presence Questionnaire categorized as questions exploring the control factor. We have not included more questions in the realism cluster as most of these referred to multimodal sensory perception and/or included sound, which was not used in our experiments. Further, it can also be noted that nine of these ten questions exhibit a strong correlation between their individual score and the total PQ score (actually among the highest correlations for all questions in the PQ).¹ We used the virtual environment described above to stage different types of causal situations. We posited that causal perception can be triggered by automatically generating co-occurring events, the first event consisting of a collision triggered by the user, provided there exist no temporal gaps between those events. Our control group thus corresponds to an environment where no such co-occurrences are created. Namely, when a subject drops a glass, the glass floats in midair before eventually landing on the table after a few seconds, without that landing being followed by any specific consequence. This behavior introduces temporal gaps (which are known to impair causal perceptions) while also possibly decreasing perceived motion transfer (through slow or irregular motion of the falling object). For this control group, the causal engine was not activated; rather, scripted behaviors were randomly selected, each corresponding to the select object returning to its original place unaltered (Figure 7).

In a first group, subjects will be presented with realistic event co-occurrences, which should induce causal perceptions (Figure 6). A third group of subjects (Figure 8) will be presented with deliberately “unrealistic” effects, that is, behaviors not semantically or physically related to the initial action (e.g., following the impact of a glass on the table, a nearby glass’s contents would

evaporate). The presence score for these two groups is to be compared to the control group, while simultaneously assessing the actual level of causal perception in each group, through the analysis of verbal explanations provided by subjects.

5.1 Experimental Setting

A total of 53 subjects were recruited and allocated to the three groups above. Group 1 was composed of 16 subjects (average age 22.6; 8 male, 8 female), Group 2 was composed of 20 subjects (average age 27.8; 9 female and 11 male) and Group 3 was composed of 17 subjects (average age 26.9; 6 female and 11 male). The average duration of a session was 30 minutes and each subject was rewarded for participation with a £15 (equivalent to \$28) voucher.

Subjects were introduced to a desktop 3D virtual environment supporting interactions with the virtual world’s objects. The environment is represented in Figures 6–8 and is composed of five tables, each supporting two glasses (one empty, one full), a beer bottle, a lit candle, and a cardboard menu.

Subjects were facing an 18-inch screen from a distance of 30–45 cm. The corresponding field of vision in the virtual environment was approximately 80°. In addition, they operated in a silent room.

After having the basic interaction mechanisms for grasping, lifting, and dropping objects explained to them in a similar but different environment (including a short training session to familiarize themselves with the system controls), subjects were given instructions for the “task” they had to carry out. For each table in the virtual world, the task was to select the empty glass, lift it above the table, then drop it and let it fall on the table, aiming at a specific virtual marker drawn on the table (this in order to avoid unwanted situations, such as subjects dropping the glass on other objects on the table).

They would then witness the virtual world reaction to their actions; in other words, the consequences of the falling pint glass hitting the table. The subjects would interact with the virtual objects using the controls provided by the native game engine: through a combination of mouse buttons and mouse movements, they are

1. This point has been criticized by Slater (1999) in his reply to Witmer and Singer, and the reader is referred to that reply. Our emphasis in this paper is, however, on presence *factors*, in particular those dealing with control and predictability.




Figure 6. Example of co-occurrence generated by the system with a high level of plausibility (experiment group 1).




Figure 8. Example of co-occurrence generated by the system with a low level of plausibility (experiment group 3).




Figure 7. Example of an “absence of causality” scenario (control group).

able to select and move objects in the 3D world. Visual feedback was provided for object selection, as well as object position above the table (through a virtual shadow, disappearing when a sufficient height had been reached, signaling the object could be dropped).

After each interaction, the subjects were asked to give a short textual explanation of the observed events, which they entered directly on the computer used for the experiments. The rationale is that explanations, rather than simple descriptions, would force the expression of causal concepts relating their actions to the observed system response. These explanations were to be used in analyzing whether subjects actually attributed causality between the events they observed. We thus collected four short textual explanations for each subject taking part in these experiments (i.e., a total of 64 for Group 1, 80 for Group 2, and 68 for Group 3). The average length of one textual explanation was 20 words for Group 1 and 30 words for both Group 2 and Group 3. The goal was to analyze these answers for the occurrence of causal explanations, hence validating the existence of causal perception in any given experiment. The analysis of free text explanations was also a way to determine how implausible events were perceived or judged, and whether mechanistic explanations were invented for them.

After completing their participation in the experiments, subjects were asked to fill in a questionnaire

composed of the ten items extracted from the Presence Questionnaire (Witmer & Singer, 1998). The questionnaire was presented in a paper form and subjects had to respond by putting an X on the continuous 7-grade scale (Figure 9).

Each of these ten questions (shown in Figure 10) was presented to the subjects as a continuous seven-point scale (from 0 to 6), where extremities, as well as the middle point, were associated textual descriptions (as originally described in Witmer and Singer, 1998). Care was taken, at all times, not to mention the words “cause” or “causality” in the experiments’ instructions, in the questionnaire elements, or in the textual descriptions underlying the grading scale.

5.2 Results and Discussion

The presence scores for each group were calculated by adding the responses to all 10 selected questions (Figure 11) on their 0–6 scales. Figure 12 represents the presence scores for each of the three groups with their error margins. For Group 1 (the group with realistic causal effects) the average presence score was 46.28; for Group 3 (the group with unrealistic causal effects) it was 33.92, and for the reference group, Group 2 (absence of causality) it was 30.82.

Q1- How responsive was the environment to actions that you initiated (or performed)?




Figure 9. Example question and its associated response grade scale.

Q1- How responsive was the environment to actions that you initiated (or performed)?

Q2- How natural did your interactions with the environment seem?

Q3- How much did your experiences in the virtual environment seem consistent with your real-world experiences?

Q4- Were you able to anticipate what would happen next in response to the actions that you performed?

Q5- How much delay did you experience between your actions and expected outcomes?

Q6- How well could you move or manipulate objects in the virtual environment?

Q7- How much did the control devices interfere with the performance of assigned tasks or with other activities?

Q8- How well could you concentrate on the assigned tasks or required activities rather than on the mechanisms used to perform those tasks or activities?

Q9- How much were you able to control events?

Q10- How quickly did you adjust to the virtual environment experience?

Figure 10. Questionnaire.

An ANOVA computed over the presence scores revealed a significant effect of group, $F(2, 50) = 20.17$, $p < .001$. Tukey HSD post-hoc tests confirmed that presence scores for the realistic causality group were higher than those in the absence of causality and unrealistic causality groups ($ps < .001$), and that these latter two groups did not differ. In each group, subjects were asked to enter on-screen a brief explanation of the phenomena observed after

each trial. They were specifically instructed to explain what happened, rather than just describe the events they had witnessed. The set of explanations (five for each subject) was pooled over individual groups and was subsequently analyzed for causal explanations. One problem in the interpretation of these textual explanations is, of course, the use of language. Although sometimes a simple juxtaposition of descriptions can constitute an implicit causal statement (see, e.g., Oes-




Figure 11. Scores obtained per question/per group (1-2-3).

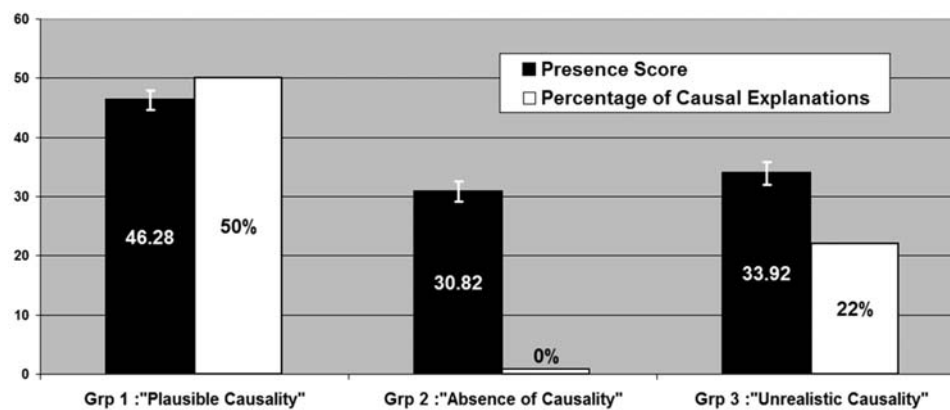


Figure 12. Presence score and production of causal explanations for each group.

termeier & Hesse, 2001), we could only interpret descriptions making explicit use of causal vocabulary.

The method we decided to use consists of analyzing each individual explanation for causal expressions corresponding to linguistic descriptions identified by Wolff (2003). We have retained the following expressions as causal:

- Explicit causal vocabulary (causes, causing, caused by), as in, *the glass fell and smashed causing the bottle to fall over, glass shattered causing candle to fall over*

- Lexical causatives (verbs that allow speakers to describe a causal situation in a single clause, as listed in Wolff (2003), for example, *when dropping the glass it moved the other glass along the table.*
- Two-argument activity verbs (also listed in Wolff, 2003) whenever their effects are also mentioned (to overcome one of Wolff's objections), as in the following, *glass shattered also knocking card over or when dropping the glass, it broke and the pieces hit the candle which in turn fell over.*

It can be noted that such vocabulary encompasses both the reporting of causal perception and the production of more sophisticated mechanistic explanations. Our method consists of considering each individual explanation, and considers it as a causal explanation if it contains one or more of the above causal expressions. We then compute the ratio of causal explanations for the whole group of subjects. Figure 12 shows the results of that analysis. For the reference group (Group 2), there were no detectable causal explanations (0 percent score). The few occurrences of causal (or mechanistic) vocabulary (e.g., *because*) were not referring to the events observed but rather to the subject's own analysis of performance (e.g., . . . *which missed the mark because I had moved the mouse after the glass began to drop*). For Group 1 (plausible causality), the level of causal explanations was approximately 50 percent. For Group 3 (unrealistic causality) it was 22 percent. Overall, a significant number of subjects simply neglected to give the explanations they were asked for, and merely gave descriptions of the events without any such explanation. This obviously affected the absolute number of causal explanations, although in a uniform fashion across groups.

The detailed proportions of causal explanations provided by participants were 32/64 for the realistic causality group, 0/80 for the non-causal group, and 15/68 for the unrealistic causality group. A Chi-Square test revealed that the frequencies of causal explanations were distributed differently between the three groups, $\chi^2(2) = 51.52, p < .001$. Because the proportion of causal explanations was 0 in group 2 ("absence of causality"), we also performed another Chi-Square test on the realistic and unrealistic causality groups only. This test also revealed a significant difference in distribution of causal explanations between those groups, $\chi^2(1) = 11.23, p < .001$. Some of the causal explanations provided by the subjects are shown in Figure 13. Several subjects perceived a causal link between co-occurring events, but provided, in addition, mechanistic explanations, such as the fact that "vibrations" accounted for the perceived causality (as in *the vibrations due to the fall of the pint have broken the bottle*). This is consistent with reports linking causal perceptions to mechanistic explanations (Schlottmann, 1999); although with this analysis we

were not able to observe any instance of dissociation between causal perception and causality judgment (Schlottmann & Shanks, 1992).

The presence score was significantly higher for Group 1 (plausible causality) than for the control (Group 2), with a difference in total score of approximately 50% with respect to the control group. Confirmation of causal perception in Group 1 comes from the level of causal explanations in this group, while confirmation that Group 2 indeed behaved as a reference group can be derived from the absence of causal descriptions.

Secondly, no major difference was observed between Group 3, in which the system produced highly unrealistic associations, and Group 2, the reference group where effects were selected not to elicit causal perception (even though the difference observed was found to be statistically significant). There are several possible explanations for this observation. The first one would consider that "realism" contributes most significantly to the presence score, even with our specific selection of *control* questions (some of which are categorized by Witmer and Singer (1998) as involving both aspects, and some also occur in the "Reality Questionnaire" under a slightly rephrased form.² In that sense, the unrealistic behavior observed in Group 3 would be less likely to produce high presence scores with the use of PQ. On the other hand, the *control* factor of presence, also being defined in terms of anticipation (Witmer & Singer), would naturally be affected by the occurrence of unrealistic effects.

But the simultaneous analysis of the verbal explanations suggests another explanation, due to the low level of causal perception in Group 3, at 22 percent, which is that some effects in Group 3 actually failed to induce causal perception. This is further confirmed by the occurrence, in the textual explanations of Group 3, of explicit statements of surprise or incomprehension (e.g., *I have no explanations for that, it surprised me because . . .*). The mixed results observed for Group 3 could be ex-

2. The importance given to realism could constitute a limitation of the Presence Questionnaire, as it would rule out presence in some purposefully unrealistic environments (artistic installations or fantasy/narrative worlds).

Group 1 “Realistic Causality”

Glass shattered causing candle to fall over.

When dropping the glass it moved the other glass along the table.

Glass shattered also knocking card over.

The glass broke and it moved the other glass to the corner of table.

The pint smashed which knocked the menu over.

The fractured glass pieces caused the menu to fall over.

When dropping the glass it broke and the pieces hit the candle which in turn fell over.

It smashed into lots of pieces and knocked over the beer bottle.

The glass smashed and knocked over the bottle.

The glass fell and smashed causing the bottle to fall over.

Group 3 “Unrealistic Causality”

The glass fell and broke the beer bottle next to it.

The pint full of beer broke. It happened because it was very close of the point of impact of the pint.

The glass fell and appeared to knock a candle off the table even though the candle did not seem to be in the path of the glass.

The bottle of wine broke. I think it happened because it was the only closed recipient on the table and the vibrations, due to the fall of the pint, have broken the bottle.

Comments expressing surprise in Group 3:

The beer in the pint disappeared like if somebody was drinking it. I have no explications for that because it could not happen in real life.

Now the cross is on the left and when I released the glass the bottle broke I did not expect it. It surprised me because I was trying to put the glass on the cross.

I was very surprised this time, the glass which contains beer did not get empty but it exploded when I released the glass. I expected that something happened but not a explosion.

When the glass was dropped the bottle in the background split. The top half of the bottle landed on the table next to the bottom half of the bottle and the liquid from bottle spill on the table.

Figure 13. Examples of causal explanations provided by subjects for Group 1 and 3.

plained by the fact that some “unrealistic events” appear more unrealistic than others.

6 Conclusions

The specific contribution of this work is to attempt to relate one well-observed psychological phenomenon, causal perception, to some fundamental ideas of presence, namely the action-based conception of presence (see, e.g., Zahorik & Jenison, 1998; Held & Durlach, 1992; Sheridan, 1992). Throughout this paper, we have referred to “conceptual” determinants of

presence as originally introduced by Witmer and Singer (1998), although they lack the validation of factor analysis (Schubert, Friedmann, & Regenbrecht, 2001). It is interesting to note, however, that in their paper on the analysis of the respective contribution of presence factors, Schubert et al. (2001) have only attributed a minor role to “control and predictability” (stating that it would account for only 2.9% of variance in presence scores), making it one of the least significant determinants of presence. In these experiments, causal perception appears to be correlated with higher presence scores, a finding that cannot be totally accounted for by realism factors, as some of the co-occurrences gen-

erated already departed from strict physical realism (hence our use of the term *plausible causality*). One question that we could not answer completely concerns the exact impact of unrealistic cause-effect associations on presence. In that experiment, subjects produced a lower proportion of causal explanations and sometimes clearly stated their disbelief at some observed effects. It could indeed be the case that some “implausible” generated effects could actually violate certain principles of causal perception, such as feature transfer or motion amplification. This should probably be revisited after gaining a better understanding of causal perception in realistic environments, which will include knowledge of relevant perceptual features inducing causal perception.

Overall, our findings suggest that causal perception could play a significant role in VR experiences, which has several implications. The first one is that causal perception could become a topic of investigation and measurement in presence studies, which would benefit from specific methodological work developed in cognitive psychology. The second one is that the induction of causal perception could become part of VR technology, offering the prospect of creating new experiences, in particular by departing from real-world physics. We have been exploring this approach, which we termed alternative reality, in the field of VR arts (Cavazza et al., 2005). This however raises the issue of presence in unrealistic worlds, as the type of questionnaire we have been using in this research still includes realism as one of the determinants of presence.

Acknowledgments

This work has been funded in part by the European Commission, through the ALTERNE Project (IST-38575).

References

- Buehner, M. J., & Cheng, P. W. (2005). Causal learning. In K. J. Holyoak & R. Morrison (Eds.), *Handbook of thinking & reasoning* (pp. 143–168). Cambridge, UK: Cambridge University Press.
- Buehner, M. J., & May, J. (2003). Rethinking temporal contiguity and the judgment of causality: Effects of prior knowledge, experience, and reinforcement procedure. *Quarterly Journal of Experimental Psychology*, 56A(5), 865–890.
- Cavazza, M., Lugrin, J.-L., Crooks, S., Nandi, A., Palmer, M., & Le Renard, M. (2005). Causality and virtual reality art. *Proceedings of the Fifth Conference on Creativity and Cognition*, 4–12.
- Cheng, P. W. (1993). Separating causal laws from casual facts: Pressing the limits of statistical relevance. In D. L. Medin (Ed.), *The psychology of learning and motivation: Advances in research and theory* (vol. 30, pp. 215–264). San Diego, CA: Academic Press.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104(2), 367–405.
- Collins, J., Hall, N., & Paul, L. A. (2004). *Causation and counterfactuals*. Cambridge, UK: MIT Press.
- D’Espagnat, B. (2002). *Traite de physique et de philosophie*. Paris: Fayard (in French).
- Fikes, R. E., & Nilsson, N. J. (1971). STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2(3–4), 189–208.
- Galavotti, M. C. (2001). *Causality, mechanisms and manipulation*. Available at <http://philsci-archive.pitt.edu/archive/00000132/>.
- Held, R., & Durlach, N. (1992). Telepresence. *Presence: Teleoperators and Virtual Environments*, 1(1), 109–112.
- Jiang, H., Kessler, G. D., & Nonnemaker, J. (2002). DEMIS: A dynamic event model for interactive systems. *ACM Virtual Reality Software Technology 2002*.
- Kruschke, J. K., & Fragassi, M. (1996). The perception of causality: Feature binding in interacting objects. *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society* (pp. 441–446). Hillsdale, NJ: Erlbaum.
- Lewis, M., & Jacobson, J. (2002). Games engines in scientific research. *Communications of ACM*, 45(1), 27–31.
- Loomis, J. M. (1992). Distal attribution and presence. *Presence: Teleoperators and Virtual Environments*, 1(1), 113–119.
- Mantovani, G., & Riva, G. (1999). “Real” presence: How different ontologies generate different criteria for presence. *Presence: Teleoperators and Virtual Environments*, 8(5), 540–550.
- McGreevy, M. W. (1992). The presence of field geologists in

- Mars-like terrain. *Presence: Teleoperators and Virtual Environments*, 1(4), 375–403.
- Michotte, A. (1963). *The perception of causality*. New York: Basic Books. Translated from the French by T. R. and E. Miles.
- Oestermeier, U., & Hesse, F. W. (2001). Singular and general causal arguments. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the 23rd Annual Conference of the Cognitive Science Society* (pp. 720–725). Mahwah, NJ: Erlbaum.
- Pearl, J. (2000). *Causality: Models, reasoning and inference*. Cambridge, UK: Cambridge University Press.
- Riva, G., Davide, F., & Ijsselstein, J. (2003). *Being there: Concepts, effects and measurements of user presence in synthetic environments*. Amsterdam: IOS Press.
- Roberts, D. J., & Sharkey, P. M. (1997). Minimising the latency induced by consistency control within a large scale multi-user distributed virtual reality system. *IEEE International Conference on Systems, Man and Cybernetics*.
- Schloerb, D. W. (1995). A quantitative measure of telepresence. *Presence: Teleoperators and Virtual Environments*, 4(1), 64–80.
- Schlottmann, A. (1999). Seeing it happen and knowing how it works: How children understand the relation between perceptual causality and knowledge of underlying mechanism. *Developmental Psychology*, 35, 503–517.
- Schlottman, A., & Shanks, D. R. (1992). Evidence for a distinction between judged and perceived causality. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 44, 321–342.
- Scholl, B. J., & Nakayama, K. (2002). Causal capture: Contextual effects on the perception of collision events. *Psychological Science*, 13(6), 493–498.
- Scholl, B. J., & Tremoulet, P. (2000). Perceptual causality and animacy. *Trends in Cognitive Sciences*, 4(8), 299–309.
- Schubert, T., Friedmann, F., & Regenbrecht, H. (2001). The experience of presence: Factor analytic insights. *Presence: Teleoperators and Virtual Environments*, 10(3), 266–281.
- Sheridan, T. B. (1992). Musings on telepresence and virtual presence. *Presence: Teleoperators and Virtual Environments*, 1(1), 120–125.
- Slater, M. (1999). Measuring presence: A response to the Witmer and Singer questionnaire. *Presence: Teleoperators and Virtual Environments*, 8(5), 560–566.
- Ware, C., Neufeld, E., & Bratram, L. (1999). Visualizing causal relations. *IEEE Visualization 1999*.
- White, P. A., & Milne, A. (1999). Impressions of enforced disintegration and bursting in the visual perception of collision events. *Journal of Experimental Psychology: General*, 128, 499–516.
- Wilkins, D. E. (1988). Causal reasoning in planning. *Computational Intelligence*, 4(4), 373–380.
- Witmer, B. G., & Singer, M. J. (1998). Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleoperators and Virtual Environments*, 7(3), 225–240.
- Wolff, P. (2003). Direct causation in the linguistic coding and individuation of causal events. *Cognition Journal*, 88, 1–48.
- Wolff, P., & Zettergren, M. (2002). A vector model of causal meaning. *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Zahorik, P., & Jenison, R. L. (1998). Presence as being-in-the-world. *Presence: Teleoperators and Virtual Environments*, 7(1), 78–89.
- Zubiri, X. (2001). *Sobre la realidad*. Madrid: Alianza Editorial (In Spanish, original work published in 1966).