A Comprehensive Taxonomy
for Three-dimensional Displays[†]

Waldir Pimenta
Departamento de Informática
Universidade do Minho
Braga, Portugal
wpimenta@di.uminho.pt

Luís Paulo Santos
Departamento de Informática
Universidade do Minho
Braga, Portugal
psantos@di.uminho.pt

ABSTRACT

Even though three-dimensional (3D) displays have been introduced in relatively recent times in the context of display technology, they have undergone a rapid evolution, to the point that a plethora of equipment able to reproduce dynamic three-dimensional scenes in real time is now becoming commonplace in the consumer market.

This paper's main contributions are (1) a clear definition of a 3D display, based on the visual depth cues supported, and (2) a hierarchical taxonomy of classes and subclasses of 3D displays, based on a set of properties that allows an unambiguous and systematic classification scheme for three-dimensional displays.

Five main types of 3D displays are thus defined –two of those new–, aiming to provide a taxonomy that is largely backwards-compatible, but that also clarifies prior inconsistencies in the literature. This well-defined outline should also enable exploration of the 3D display space and devising of new 3D display systems.

Keywords

three-dimensional displays, depth cues, 3D vision, survey, taxonomy

1. INTRODUCTION

The human ability for abstraction, and the strong dependence on visual information in the human brain's perception of the external world, have led to the emergence of visual representations of objects, scenery and concepts, since pre-historical times. Throughout the centuries, many techniques have been developed to increase the realism of these copies.

Recent years have revealed a focusing of these efforts in devising ways to realistically recreate the sensation of depth, or three-dimensionality, of the depicted scenes. 3D displays thus emerged as an active area of research and development.

Despite this being a relatively recent field, many different approaches for 3D displays have been already proposed and implemented, and new ones surface with some regularity. Moreover, these implementations provide different sets of approximations for the depth cues that our visual system uses to perceive the three-dimensionality of a scene.

This profusion of implementations has plagued attempts to define a nomenclature system for 3D displays. While a few comprehensive classification schemes have been proposed, most based on or compatible with Okoshi's seminal work [Oko76], several are now obsolete, while recent attempts tend to be specific to a subset of displays [MK94, Hal97, BS00, Dod05, GW07, UCES11].

Mostly what is seen are overview sections in publications that on the one hand assume implicit definitions of 3D perception and 3D displays, and on the other hand frequently avoid taking a stance (or do so inconsistently) in undecided issues emerging from partially incompatible previous classifications, such as the placement of holographic technology [Fav05] or integral imaging [DM03]. A definitive, exhaustive and unambiguous categorization system for 3D displays has thus been lacking in the literature [CNH+07, p.1], which hinders the classification and evaluation of different implementations, especially hybrid ones.

The approach presented in this paper focuses in the formalization of the properties of each category of 3D displays, to provide a stable system for classifying existing or new implementations. Specifically, the definition and categorization of 3D displays is based in their fundamental properties, rather than in implementation details, as is the case with most current classifications.

As a necessary foundation for this taxonomy, Section 2 presents a general overview of the depth cues used by the human visual system to perceive three-dimensionality. With this knowledge, we can then, in Section 3, determine the specific subset of these that clearly mark the frontier between 2D and 3D displays, and define basic properties of 3D displays. Section 4 then delves into the 3D display realm, defining a hierarchy of types and subtypes for 3D displays, based primarily on the depth cues they implement.

By employing a systematic approach, we expect the outcome to be a logical, well-structured and extensible taxonomy that will facilitate comparison of different approaches, and the evaluation of appropriate techniques for a given application. The Conclusion assesses the degree to which this objective was fulfilled, and illuminates what further work is to be performed to complement the proposed taxonomy.

2. VISUAL CUES TO THREE-DIMENSIONALITY

The origins of the Human species, as primates living and moving in trees, and later as hunter-gatherers, contributed significantly to make perception of depth a very important feature of our vision. Developments in art and research in optics and display technology have revealed some of the cues that our visual system uses to interpret the location of objects. These hints, known as **depth cues**, can be divided into two main groups: *psychological* cues, which depend on acquired knowledge of the visual aspect of familiar objects, and *physiological* cues, which manifest through the anatomy of our visual system [Oko76].

The main psychological depth cues are:

Occlusion. The overlap of some objects by others that are closer to us. This is one of the most fundamental ways we perceive depth on a scene.

Linear perspective. Given prior knowledge of common shapes and/or sizes of objects, we interpret perceived distortions in their shape (parts farther away from us appear smaller), differences in size between them, and variation of their angular size (how much of our visual field they cover) as indicators to their location in three-dimensional space.

Atmospheric perspective. Commonly known as "distance fog", it refers to the fading in contrast and detail, and shift to bluish colors, of objects located at a great distance. This happens because the light we get from them had to travel an increased distance through air and thus underwent more scattering from the atmospheric gases and particles.

Shading and **shadow projection.** Effects caused by the relationship between objects and light sources. The distribution of brightness and color in an object's surface provides information about (among other things) its shape and position relative to the light sources that illuminate it. Also, the location, format and darkness of shadows projected into the object (due to parts of it or other objects obscuring the light) and into its vicinity allow us to interpret its 3D form and relative position to other objects and/or the environment.

The above are all static cues. There are two more psychological cues, which are dynamic; that is, they manifest when there is movement either of the observer or of the observed object (or both):

Motion parallax. Relative changes in perceived position between two objects when we move. For example, during a car trip a tree seems to be "travelling past us" faster than the distant mountains.

The kinetic depth effect. Changes in the appearance of an object due to its own motion. For example, when a spherical object –say, a football– is uniformly illuminated so that no shadows give away its round shape, a slow rotation around itself is sufficient for our visual system to infer that it is a solid body and not a flat disk facing us, due to the relative motions of features in its surface.

The physiological depth cues consist of:

Binocular disparity (or **stereo parallax**)[1]. Differences in images received by each eye, commonly called **stereoscopy**[2]. Studies indicate [Oko76] that for a moderate viewing distance, binocular disparity is the dominant depth cue to produce depth sensation, through a process called *stereopsis*, which is the effort made by the brain to fuse the images together into a 3D perception of the scene. This fusion effort is always necessary because convergence of the eyes can only produce a perfect match for a limited subset of the points from the images, due to projection geometry constraints.

Convergence. When both eyes rotate inwards to aim at the object of interest, thus aligning the different images they receive, so they can be more effectively combined by the brain. As with accommodation, this rotation manifests itself with greater amplitude when differences in distance occur closer to the eye, so it is also a cue that is more strongly perceived for nearby objects (less

---

[1]"Binocular" comes from the Latin *bini* (pair) + *oculus* (eye). "Stereo" comes from the Greek *stereós* (solid).

[2]It's been known since as early as 300 B.C. that depth perception in human vision is related to the fact that we have two eyes, in separate physical locations, which collect different simultaneous perspectives of the same object [EucBC].

than 10m, according to [Wid01]). If they are close enough, one can clearly feel the eyes "crossing" so that they can keep aiming at the same point.

Accommodation. The effort made by the muscles in the eye that control the shape of its lens in order to bring the image into focus in the retina. Even though we usually do not consciously control these actions, our brain uses this muscular contraction information as an indicator of the distance of objects we are observing. Since the focusing effort varies much more for distance changes near the eye, the effect is particularly notable for nearby objects (less than 2m, according to [MZ92]).

The depth cues described above are summarized in Table 1.

The Accommodation-Convergence Mismatch

The fact that most visual representational media are unable to implement all depth cues –especially the physiological ones–, does not pose a serious problem, either because the scenes represented are meant to take place (or be viewed from) a distance where the physiological cues aren't relevant [Oko76, p.39], or because we can cognitively ignore the mismatch in psychological vs. physiological depth cues, as our abstraction ability allows us to understand their purported three-dimensionality regardless.

However, a mismatch *among* the physiological cues is less tolerable. This mismatch is common in current 3D displays, because every display that provides stereoscopy (one view for each eye) is theoretically able to implement proper convergence cues for each object in the scene depending on their location. But accommodation (provided by the ability to make the light rays diverge not from the screen, but from the virtual positions of the scene objects) is much harder to achieve; therefore, most of these displays end up forcing the eye to always focus at the screen to get a sharp image, which conflicts with the cues of convergence and stereopsis. The resulting phenomenon is called the *accommodation-convergence mismatch*.

This mismatch is more serious than the aforementioned one, because providing the brain with conflicting physical signals causes discomfort, the same way mismatch between visual and vestibular (from the balance system in the inner ear) perception of movement causes motion sickness. The consequences may include headaches, fatigue or disequilibrium, preventing continued use of these displays. This, of course, in addition to the reduction it causes in the realism of the 3D visualization, which might become uninteresting or even visually confusing [Hal97].

Table 1: Summary of visual depth cues for three-dimensional vision

	static	dynamic
psycho-logical	occlusion (overlap); linear perspective; atmospheric perspective (distance fog); shading and shadows.	motion parallax; kinetic depth effect.
physio-logical	accommodation (focus); binocular disparity (stereoscopy); convergence.	

3. DEFINITION OF A 3D DISPLAY

Before defining what a 3D display is, it is necessary to clarify what is meant by "display". As a word with multiple meanings, we will assume the context of visual perception and the word's usage as a concrete noun (i.e., the name of a thing). As such, the definition adopted will be "a visual output device for the presentation of images".

It's worth pointing out that the word "images" is in plural, because we will consider only display media that don't produce permanent records, but instead are mutable, or rewritable, by comprising reconfigurable active elements, such as pixels, voxels[3] or catoms[4] – in other words, *electronic visual displays*. This, as [Oko76] pointed out, effectively excludes static visual representations such as paintings, photographs, sculptures, and even classical (static) holograms, for they are not displays in the sense adopted above, but merely the physical embodiment of a specific image. These will therefore be left out of this taxonomy. Nevertheless, all the principles behind them are present in the displays we consider, the only difference being the adoption of a rewritable medium.[5]

With the clarification of what constitutes a display device, we can now approach the question of what makes a display three-dimensional. Firstly, we must acknowledge that the line separating 3D displays from 2D displays is not always clearly defined, despite what the dichotomic "2D/3D" nomenclature seems to suggest. This fuzziness occurs because, on the one hand, the psychological 3D depth cues can, in fact, be reproduced in media traditionally considered as 2D; and on the other hand, many displays deemed three-dimensional are actually flat screens, which means that the images are emitted from a two-dimensional surface.

---

[3] A *portmanteau* of the expression "volumetric pixels".
[4] In the (still theoretical) field of claytronics –dynamic sculptures made of microscopic robots–, "catom" is a combination of the words "claytronic atoms" [GCM05].
[5] For instance, when we mention holography, or stereoscopic displays, we will be referring to their electronic counterparts.

With these limitations in mind, **we define 3D displays as visual output devices that evoke at least one of the physiological depth cues** (stereoscopy, accommodation and convergence) – besides, naturally, the psychological cues enabled by the specific display technology used. This definition ensures that the 3D perception is truly engaged in a natural way, and not by ignoring the apparent flatness of the scene, as happens with displays based only in psychological depth cues.

4. PROPOSED TAXONOMY FOR 3D IMAGING TECHNIQUES

To define a basis for the proposed taxonomy, we will apply two general criteria as orthogonal axes of categorization. We'll demonstrate that by intersecting these two basic properties, it is possible to establish a well-grounded, formally-defined taxonomy that largely validates current consensus but also clarifies conflicting definitions.

The first axis is the **number of views supported** by the display. The reasoning behind this is that most of the depth cues for 3D perception (occlusion, motion parallax, convergence, stereopsis, etc.) are dependent on the angle from which the observer views the scene. 3D displays will employ different methods to emulate this viewpoint-dependent variation of the light field.

One such method consists simply in producing two views and ensuring that each is only seen by the appropriate eye of the observer. Another approach employs displays that are able to project multiple views into different directions. This is implemented by segmenting the image into as many perspectives as desired, multiplexing them into the display, and using a filtering mechanism to direct each view to the corresponding direction. Finally, a third type comprises displays that can generate or approximate a continuous wavefront of light that propagates as coming from the actual 3D position of the virtual object, rather than dispersing from its projection in the display surface.[6]

Throughout the years, as 3D displays advanced past the two-views (binocular) approach, the word "stereoscopic" has gradually expanded its range to become largely synonymous with three-dimensional vision (and rightly so), and is thus routinely applied to displays of all of these types. Therefore, in the spirit of unambiguity, the three meta-categories described above will be named *"duoscopic"*, *"multiscopic"* and *"omniscopic"*, respectively.[7]

The other main axis we'll use to map the 3D displays space is the **effective shape of the display medium** itself, which can be *"flat"* or *"deep"*. This doesn't depend strictly on the shape of the display surface, but rather on the effective volume it occupies while displaying the 3D image. The flat displays can be compared to a window, a planar surface which provides different perspectives as one moves around, but limits the scene at its boundaries. For the deep displays, there is a volume of space occupied by the display medium (either permanently or due to moving elements) and the virtual object is displayed inside the volume, also not able to exist outside the volume's boundaries as they are perceived by the observer. We can say that one looks *through* flat displays as if through a window, and looks *into* deep displays as if they were a crystal ball.

Aside: the projection constraint

The boundary limitation of both the flat and the deep displays are manifestations of the *"projection constraint"*.[8] Countering this effect may be done by increasing the absolute size of the display (for example, a cinema screen), shaping it in order to surround the viewer (as is done in the CAVE virtual reality environment), or increasing its relative size by bringing it closer to the observer (the technique used by virtual reality glasses).

These two criteria allow us to effectively separate the displays into five main categories, most of which are already well-established in the literature. Table 2 summarizes this division.

It might be noticed that two of those terms are not common in most taxonomies, namely "virtual volume displays" and "multi-directional displays". They are, in fact, key components of this taxonomy, in that they clarify the classification of techniques for which past works have not been able to agree on a category. Other categories, however, were included with their currently *de facto* standard names, in order to prevent excessive disruption and preserve as much backwards-compatibility as sustainable without breaking the consistency of the proposed framework.

Table 2: Proposed Taxonomy

		display shape	
		flat	*deep*
# views	*duoscopic*	stereoscopic	
	multiscopic	autostereoscopic	multi-directional
	omniscopic	virtual volume	volumetric

---

[6]Head tracking by itself only implements monocular directional variation; thus, it doesn't constitute a 3D display as defined above.

[7]Prior attempts to define the difference between these types of displays have entailed the use of the terms "stereograms" and "panoramagrams" [Oko76, Hal97], but the distinction hasn't been widely adopted in the literature, and even less in the industry.

[8][Hal97] describes the projection constraint by stating that "a display medium or element must exist in the line of sight between the viewer and all parts of the [visible] image."

In the following subsections we will complete the definition of these five groups by specifying their main properties and, where applicable, defining relevant subcategories inside them.

Flat 3D displays

Flat-type, screen-based 3D displays are the most popular kind of 3D displays used currently, with commercial use now common in movie theaters and domestic entertainment devices. They work mostly by providing stereoscopy (different images for each eye), which, as mentioned in Section 2, is the main depth cue for 3D vision at moderate distances.

These displays can be further divided in three main groups: **stereoscopic** devices, which work in conjunction with glasses to provide two distinct views; **autostereoscopic** screens, which can generate multiple views without requiring any headgear; and **virtual volume** displays, which recreate the 3D wavefront as if propagating from the actual location of the 3D image – the most notable example being the hologram.

4.1.1 Stereoscopic Displays

Stereoscopic 3D displays can display one image to each eye in two ways: either by combining (i.e, multiplexing) two separate streams of images in one device, and filtering them with special glasses, or by using separate display devices for each eye.

Glasses-based stereoscopic displays can be implemented through three filtering techniques [Ben00]:

Wavelength multiplexing. Separating the left-eye and right-eye images in different colors, the most well-known example of which is the anaglyph, with its characteristic "red-green" glasses;

Temporal multiplexing. Using shutter glasses synchronized with the screen and a doubled frame-rate that displays the images for the left and right eye alternatively;

Polarization multiplexing. Achieved by emitting images for each eye with different light polarizations (direction of wave oscillation), and filtering them with polarized-filter glasses.

The stereoscopic displays that use separate screens for each eye are usually called **head-mounted displays (HMDs)**. This name is justified because the whole display system is head-mounted, rather than only the filtering mechanism.

HMDs include mostly devices such as virtual reality (VR) or augmented reality (AR) glasses, but also comprise techniques still largely embryonic, such as retinal projection, contact lens displays and brain-computer interfaces.

As previously mentioned, HMDs can overcome the projection constraint by displaying the image closer to the eye, thus increasing its relative size and coverage of the visual field.

There are two key characteristics of stereoscopic displays that separate them from other 3D vision techniques: (1) they require either the whole display system or the filtering mechanism to be fixed regarding the eyes, which in most cases implies some sort of headgear, thus being potentially invasive to varied degrees (ranging from light and inexpensive filtering glasses to surgery-requiring neural implants), and (2) because they only present two views, they only support a single user/perspective.[9]

Motion parallax is not natively supported by stereoscopic displays, but they can be enhanced to support it by employing head tracking [Dod05].

4.1.2 Autostereoscopic Displays

Autostereoscopic screens are usually implemented using two techniques:

Parallax barriers, which work by sequentially interlacing the images for each perspective in vertical strips, and employing a fence-like barrier that restricts the light from each strip to propagate only in its corresponding direction.

Lenticular displays, which do this filtering by using an array of lenses that direct each part of the image to the correct direction. These lenses are usually cilindrical, providing only horizontal parallax, but spherical lenslets have been proposed to overcome this limitation, resulting in what is called an "integral imaging" device.

Autostereoscopic screens exploit the fact that the eyes occupy different points in space to provide stereoscopy. In other words, they employ direction-multiplex to channel information of the left and right views into appropriate eyes [DM03].

These direction multiplexing techniques can be generalized to produce more than two views, which enables motion parallax, and consequently the ability to support multiple observers with a single display, without any headgear. However, undesired optical distortions caused by too small lenses or barriers limits the number of possible views. The motion parallax supported is thus markedly non-continuous, which reduces the realism of the 3D effect [Hal97].

_____

[9]It is possible, using HMDs, to implement multi-user applications by having each user wear their own device, and keeping all of them synchronized, but this is obviously a costly and technically challenging approach.

Anisotropic diffusers (surfaces that scatter light in very narrow horizontal directions) have been presented as a potential solution to such limits [UCES11]. It's been reported[Tak06] that with enough angular resolution, such displays could even create accomodation responses in the eye. Therefore, by sufficiently approximating (in the assigned visualization area) the continouous wavefront that a real object would create, they could be considered omniscopic instead.

4.1.3 Virtual Volume Displays

Virtual volume displays, as the name says, are able to generate the sensation of depth by placing virtual images in 3D space, without having to physically span the imaging volume [Hal97]. Since each point of the image is optically located at the correct depth, these displays are able to provide proper accommodation. This can be implemented either by adaptive optics, or through the holographic technique.

Adaptive optics employ dynamic optical systems that can change their focusing power. These can be deformable (varifocal) membrane mirrors, or "liquid" lenses, usually produced through an effect called "electrowetting". They are similar to the old illusion called *Pepper's ghost*, which consists in a semitransparent mirror that superimposes a reflection (of a real object, or a verisimilar 2D projection) over the background scene, producing a ghostly image of the object, and which still finds modern use in many theme parks and live shows.

In displays based on adaptive optics, the flexible optical element will reflect or transmit a static screen that displays a sequence of depth slices, synchronized with the curvature of the mirror or lens to place the image of the slice in the appropriate depth location. This kind of display will prevent occlusion, since the virtual slices cannot block the light from those behind it. But if a single-perspective is acceptable, such as in HMDs or single-user desktop displays, occlusion can be simulated by subtracting a depth layer from those behind it (from the perspective of the observer).

While the surface of the lens or mirror is not strictly planar, slight changes in their focal length lead to large variations in the virtual image's location [DM03]. Coupled with the window-like viewing mode they enable, this means that adaptive optics-based displays can be considered flat displays.

Holography, on the other hand, works by storing the shape of the wavefront of the light emanating from the scene, by recording the interference pattern of its interaction with a clean, coherent light source. The original wavefront can then be reconstructed by illuminat-

ing the pattern with a copy of the reference coherent beam. All optical effects such as shadows, reflections and occlusions are present in the resulting image.[10]

Unlike most autostereoscopic screens, virtual volume displays can provide all the physiological depth cues (particularly accommodation), as well as continuous motion parallax. The recent advances in anisotropic screens have shortened this gap, but further properties such as vertical parallax are yet unreported in such displays, which positions virtual volume displays favorably in the realism of the 3D effect and the compactness and portability of the display system.

Deep 3D Displays

Deep displays physically occupy a volume of space and display the object inside it. Two methods can be used to implement such a system: **volumetric displays**, which place the virtual points of the object in physical 3D space, and **multi-directional screens**, which, as the name says, have either a single rotating screen, or multiple static screens facing different directions – in either case, users in a given position will see only the appropriate perspective.

Volumetric displays are omniscopic, since having the object displayed in actual 3D space allows virtually any viewpoint to get the correct perspective. Multidirectional screens will have to subdivide the perspectives into a finite number of views, and are therefore part of the multiscopic meta-category. Both can potentially implement a 360° viewing angle.

4.2.1 Volumetric Displays

Volumetric displays use several techniques to display an image in real 3D space. This means that each point of the image is actually located at the position they seem to be. This can be achieved by two main methods: static volume displays, and swept-volume displays.

Static volume displays use a substrate (solid, liquid, or gas) that is transparent in its resting state, but becomes luminous, or opaque, when excited with some form of energy. If specific points can be selectively addressed inside a volume of space filled with such a material, the activation of these points (called volumetric pixels, or voxels) forms a virtual image within the limits of the display.

Naturally, gaseous substrates are preferred, and displays have been made using artificial haze to produce unobtrusive, homogeneous clouds suspended in the air

---

[10]Holograms store the entirety of the information from a scene – hence their name, which derives from the Greek "holo", the same root that the word "whole" came from.

that make light beams visible. Purely air-based displays have also been proposed, using infrared laser light to produce excited plasma from the gases in the air, at the focal points of the laser. Advanced forms of such displays are common in science fiction, often mistakenly referred to as "holograms" [Hal97]. However, the actual visual quality of such displays is very far from their imagined counterparts, and even quite low compared to other current methods of 3D vision.

Swept-volume displays use a two-dimensional surface that cyclically sweeps through a volume (either moving from one extremity to another, or rotating around an axis) and display, at each point of this path, the corresponding slice of the virtual object. Due to the temporal persistence of vision, this results in what resembles a 3D object.

The main problem with volumetric displays is that, since most of the substrates used become bright when excited, rather than opaque, each point of the virtual object won't block light from the other points [Fav05], which undermines the very basic depth cue of occlusion; that is, observers would see the back side of objects as well as their front side. This is the same problem that plagues varifocal mirror displays. Such devices are therefore better-suited to display hollow or naturally semi-transparent objects, or non-photorealistic scenes – for example, icons, or wireframe 3D models [Hal97].

This difficulty could be surpassed in static-volume displays, if the substrate can be made opaque; however, a solid, static substrate would make direct manipulation and interaction with the object impossible (which is also true of swept-volume displays). The ideal volumetric display would thus be a "dynamic sculpture" that is able to change its shape and appearance according to the desired properties of the object being visualized. This has already been proposed, in a concept called "claytronics" [GCM05], but remains a strictly theoretical possibility, with no practical implementations produced so far.

4.2.2 Multi-Directional Displays

Recently, some claims have been made in the literature that the lack of occlusion in volumetric displays is not an intrinsic characteristic of the category, but a technical limitation that can be addressed.

While, as described above, this is true of static-volume volumetric displays, swept-volume displays are strictly unable to overcome this property because they work through persistence of vision, and therefore even if the active elements could be made opaque, no part of the image is permanently located in its physical position, so light would still pass through that space in the fractions of time where the display surface isn't sweeping through that particular location.

Still, swept-volume displays purported as "occlusion-capable" have been presented in recent research (for instance, [CNH⁺07]). They work by employing highly anisotropic diffusers to ensure that light produced or projected in the display surface is only emitted in roughly the direction the display is facing, thus ensuring that only the correct view is observed in each direction. By correctly varying the image presented in the screen according to the direction it is facing, a 3D image is produced, which can also appear to float outside the display volume.

This kind of display, however, while very similar to swept-volume *volumetric* displays, is not volumetric itself, since the image points are not located in the actual position they appear to be; in other words, they manifest the property we earlier associated with multiscopic displays, that light from each point disperses from the screen itself rather than from the correct location of the virtual point, which disables the provision of the accommodation depth cue.

These **rotating screen displays** are fundamentally similar to an earlier technique known as **cylindrical hologram** [FBS86], in which a series of images taken of a subject, with a camera performing a 360° orbit around it, are recorded in thin vertical holographic strips, which are then assembled in a cylindrical shape to provide full panoramic view of the 3D object.

In both cases, the viewer-depended variation is implemented explicitly through segmenting the viewing field, rather than producing the appropriate wavefront of the 3D scene. Cylindrical holograms, however, can potentially implement accommodation if the strips aren't holograms of a flat photograph, but of the actual 3D object itself.

5. CONCLUSIONS AND FUTURE WORK

3D displays are increasingly popular choices to provide new, more immersive and intuitive tools for education, entertainment (especially in gaming, television and cinema), telepresence, advertising, among others.

Moreover, as the technology advances, more demanding uses of such displays have started becoming feasible or expectable in the near future. Such uses require high-fidelity 3D reproductions of objects, and include areas as diverse as product design, medical imaging and telemedicine, 3D cartography, scientific

visualization, industrial prototyping, remote resource exploration, professional training and architecture.

Such wide appeal has led to the rapid development of many techniques for 3D visualization, and sometimes this has resulted in poorly-defined boundaries between techniques – especially hybrid ones. This work presented a comprehensive taxonomy of 3D displays, focusing on fundamental characteristics rather than implementation details. This property should make the taxonomy robust and expansible to include new techniques and innovations. It also provides a high-level overview of the 3D displays landscape, a useful tool for researchers entering the field.

An important property of the proposed taxonomy is that it equips both researchers and practitioners with a well-defined field map which enables application-based exploration of the 3D display space. Logically separated groups of technologies allow a faster analysis of desired properties, such as the ability to perform **direct manipulation** on the virtual objects at their apparent locations, or to **overlay** the images onto the real world, to provide augmented reality, or to operate without **headgear**. Proper **accommodation** might be crucial for high-precision applications, while support for **multiple users** is relevant in design contexts.

Furthermore, a well-defined taxonomy should also enable informed speculation over the 3D display space henceforth outlined, regarding possible new techniques and analysis of their feasibility and properties, or alternatively, discarding of a specific combination of properties (or set thereof) due to economic, physical or technological limitations. This is expected to enable new 3D display systems to be conceived. As an example, one could easily conceive a static volumetric display that provides occlusion, by using a substrate that becomes opaque when excited. This could be a relevant research topic in materials science.

This study now calls for further developments in the form of an exhaustive listing of implementations and their calatoguing in a table or database that will allow manual or automatic filtering and comparison of different display technologies and respective features.

6. REFERENCES

[Ben00]   S.A. Benton. *Selected papers on three-dimensional displays*. SPIE, 2000.

[BS00]   B. Blundell and A. Schwarz. *Volumetric Three-Dimensional Display Systems*. Wiley, 2000.

[CNH+07] O.S. Cossairt, J. Napoli, S.L. Hill, R.K. Dorval, and G.E. Favalora. Occlusion-capable volumetric 3D display. *Applied Optics*, 46(8):1244–1250, 2007.

[DM03]   S. Dudnikov and Y. Melnikov. Review of technologies for 3D acquisition and display. Technical report, EU Project IST-2001-38862 TDIS, 2003.

[Dod05]   N.A. Dodgson. Autostereoscopic 3D displays. *Computer*, 38(8):31–36, 2005.

[EucBC]   Euclid. *Optics*. 300 B.C.

[Fav05]   G.E. Favalora. Volumetric 3D displays and application infrastructure. *Computer*, 38(8):37–44, 2005.

[FBS86]   D.S. Falk, D.R. Brill, and D.G. Stork. *Seeing the Light*. chapter 14: Holography, pages 389–391. Wiley, 1986.

[GCM05]   S.C. Goldstein, J.D. Campbell, and T.C. Mowry. Programmable matter. *Computer*, 38(6):99–101, 2005.

[GW07]   T. Grossman and D. Wigdor. A taxonomy of 3D on the tabletop. In *IEEE Int'l Workshop on Horiz. Interactive Hum.-Comp. Systems*, pages 137–144, 2007.

[Hal97]   M. Halle. Autostereoscopic displays and computer graphics. *ACM SIGGRAPH Computer Graphics*, 31(2):58–62, 1997.

[MK94]   P. Milgram and F. Kishino. A taxonomy of mixed reality visual displays. *IEICE Transactions on Information and Systems*, 77(12):1321–1329, 1994.

[MZ92]   M. McKenna and D. Zeltzer. Three dimensional visual display systems for virtual environments. *Presence*, 1(4):421–458, 1992.

[Oko76]   T. Okoshi. *Three-Dimensional Imaging Techniques*. Academic Press, 1976.

[Tak06]   Y. Takaki. High-density directional display for generating natural three-dimensional images. *Proceedings of the IEEE*, 94(3):654–663, 2006.

[UCES11] H. Urey, K. V. Chellappan, E. Erden, and P. Surman. State of the art in stereoscopic and autostereoscopic displays. *Proceedings of the IEEE*, 99(4):540–555, 2011.

[Wid01]   T. Widjanarko. Brief survey on three-dimensional displays. *Media*, pages 1–27, 2001.