"What's up Prof?"
current challenges in post-production and visual effects

v3.2

Neil Dodgson, University of Cambridge
John Patterson, University of Glasgow

Revised 12-3-10

Summary

In December 2008, a team of academics visited visual effects and post-production houses in London. Our aim was to discover what challenges they face and how the universities can help. Our summary report was published in the February 2010 issue of Leonardo[1]. This, longer, report expands on that summary and adds new thoughts that have emerged from subsequent discussions with our contacts on five topics: novel interfaces (1(i)), computer vision (1(j)), intellectual property (2(i)), data volumes (2(j)), and the importance of the industry (end note).

Context

The "What's up Prof?" team visited a number of London visual effects and post-production houses over two days (3–4 December 2008) to gather information on current challenges in the industry. The team was Dr Neil Dodgson (Cambridge), Prof. Phil Willis (Bath), and Dr John Patterson (Glasgow). Prof. Peter Robinson (Cambridge) joined us for one visit.

For historical reasons most of the film effects companies are grouped in Soho in central London. A few of companies (which we did not visit) are based around the major UK film studios still functioning as such (e.g., Pinewood and Shepperton) but this is because they need to do their work on set (e.g., animatronics, explosions) rather than in the quiet of a darkened room which could be anywhere. So Soho was where we went.

We visited five companies and also met with Peter Stansfield, who is part of a consultancy (Wavecrest Systems Ltd) that advises the industry on trends and funding. The six companies represent a range of specialities within the industry. We have subsequently shown drafts of this report to other companies and taken their comments on board.

As academic researchers, the "What's up Prof?" team are interested in technologically oriented research in support of the media industries, especially movies. We are also concerned with wider application of the technologies identified, such as using drawing beautification technology for therapeutic purposes. We have met several times over the years. Before our London visit we met in Cambridge (6–7 November 2008) to identify areas in which we were interested in developing new research projects, such as would be funded by the EPSRC.

---

[1] "What's up Prof? Current issues in the visual effects and post-production industry", N. A. Dodgson, J. Patterson, and P. J. Willis, *Leonardo* **43**(1):92-93, Feb 2010, ISSN 0024-094X

Our London visit then had two aims: principally to ascertain what challenges the industry is facing, and secondarily to see what support the industry would give to the ideas behind our own research proposals. We use a listening protocol: we wanted to hear first what the industry thought were the problems they would encounter in the next decade and then later see whether and where these fitted with what we wanted to do. We did not try to present our agenda, but we responded when the industry agenda chimed with ours. We went to listen and summarise, not to promote. The main part of this report contains a summary of what industry told us.

A trip around Soho

The companies

Smoke & Mirrors undertakes a full range of post-production effects, working in advertising, television, and movies. It has over 100 employees, most of whom are effects artists.

Framestore produces visual effects for movies and, more recently, has started working on computer graphics feature animation. It has over 400 employees, of whom around 20 are technical staff and around 400 are artists. Framestore is one of the largest effects houses in Europe. When it was known as by its previous name, **Framestore CFC**, its CFC subsidiary developed its own scanners. Today, however, they use Filmlight scanners (see below).

The Foundry is a software developer. It has around 50 employees, mostly programmers. It produces software for post-production and visual effects. It sells to companies like Smoke & Mirrors and Framestore. The Foundry has the development rights to Nuke™, a compositing system formerly owned by Digital Domain. Nuke has effectively replaced Apple's product Shake, which was itself originally developed by NothingReal.

Filmlight is a systems developer. It has around 70 employees. It produces a combined software and hardware solution for colour grading, the final stage in post-production. Filmilight sells to major studios and post-production houses and has around 200 systems installed worldwide. The senior partners of Filmlight used to work at Framestore CFC, when it still had that name. Filmlight won four technical Oscars this year (2010) and two further awards at the 2010 London Export Awards.

Cinesite undertakes conversion between digital and analogue media. The company also has visual effects divisions, but our contact worked in the imaging department. Cinesite used to have its own laser film scanning and film output systems, provided by Kodak of which it is a division, but it now uses Filmlight's scanners and Arri film writers.

Wavecrest Systems Ltd is an independent consultancy, founded in 1992. Peter Stansfield, our contact, is one of the partners. Wavecrest coordinates projects across the industry.

Farblue Images was not visited on the main Soho trip, but provided input later when one of us (Patterson) met with Bill Scanlon. Farblue specialises in converting 2D imagery into 3D mostly using manual methods. They aimed to automate their data-path, but that was still for the future. It is our understanding that Farblue has gone into receivership since our meeting.

The challenges for the industry

We asked each company to discuss their current problems and desires. These fell into three categories:

1. Developing new technologies.
2. Managing infrastructure.
3. Managing people and process.

With regard to research timescales, the universities and companies differ. The companies need solutions to their current problems, on a timescale of 6 to 24 months. The universities need to work on problems that will become pressing in 5 to 10 years time or on problems for which no solution is obvious to industry. The latter are those problems to which no company will devote resources but for which a solution would be useful, if one could be found. Our principal goal was thus to identify problems that universities could work on effectively, whose solutions would be useful to this creative industry. The following three sections list the problems that were identified. Our own research groups (graphics, imaging, vision) could effectively contribute to the first of these; the other issues need to be tackled by groups in other areas or by the industry itself.

1. Desirable new technologies

a) **Human in the loop**. There is much good university research that works well at the low quality end of the market (for example, segmentation, 3-D reconstruction) or for vision problems which are not too difficult (e.g., removing an obstacle to a scenic view). However, this has had little impact on the high quality end, where everything is still done by hand. It would be useful to investigate methods that solve particular problems (for example, optical flow, boundary detection, object detection) in such a way that a human being is able either to direct the automated algorithm or to adjust the output of the automated algorithm quickly and efficiently. In either case the result should be superior to the manual method yet take less time to achieve[2]. [Filmlight, Farblue]

b) **Repurposing**. Research is needed into ways to reuse 3-D models and to reuse footage. At present, 3-D models tend to be made anew for each sequel of a movie. This is understandable as technology moves on between a movie and its sequel. However, we also find that the 3-D models used for a movie are not used for the accompanying game. How can make better use of existing assets? [Foundry, Wavecrest]

c) **Finding assets**. How do we find digital assets? How do we ensure that we find the correct version when there are many different versions? How can we catalogue digital assets? The databases of assets are now so large that we need

[2] We were told the salutary story of an academic project which produced an algorithm that would take scenes filmed in daylight and recolourise them to look like night-time scenes. The quality of the result was convincing but the calculation was computationally demanding and certainly not real-time. After a demonstration of the method, a skilled colourist was brought in. He was able to achieve the same effect with a few deft applications of a manual colourisation tool. The moral of the story was that the industry is filled with skilled operators, like the colourist, who can achieve seeming miracles quite straightforwardly by hand. Care must therefore be taken when selecting what is worth automating, as it is not obvious.

to develop better ways to search images and 3-D models. [Wavecrest, Framestore]

d) **3-D reconstruction**. Reasonable methods for the reconstruction of 3D objects exist but they work best with frame-synchronised views from binocular cameras. The next challenge is the extraction of data of good enough quality for the reconstruction of a complete 3D scene from multiple movie cameras. Some aspects of this problem remain challenging. Support for 3D (stereoscopic) movie-making has become a priority for the industry following the popularity of recent 3D releases[3]. [Foundry, Farblue]

e) **Artist-directable physical stimulation**. Movie effects need to be visually plausible but simulations do not have to be physically realistic nor work for longer than the shot. There has been considerable research on producing physically *realistic* simulation. However, there is a need to provide physically *plausible* simulation that can be directed and modified by the artist[4]. For example, could you make a water simulator where the artist can control where the water goes? Could you make a cloth simulator which is physically plausible but which gives the artist control over certain behaviours? How do we make a cloth simulator that allows the cartoon character-style squash and stretch? How do you make things that look plausible when they are physically impossible? One example given was of an animated character performing a physically impossible, high-G somersault; the high forces involved caused the cloth simulator to rip the clothes off the character. The company had to find a work-around to ensure that the clothes stayed on the character to produce the effect required by the director rather than the physically-correct effect. [Framestore]

f) **Making convincing digital humans**. Human beings are good at recognizing and analysing the appearance and behaviour of other human beings. It is difficult to make a convincing digital human. Indeed, there is evidence that a digital human that is not quite convincing is more disturbing to the average viewer than a digital human that is clearly not meant to be realistic ("the uncanny valley"[5]). Acquisition of good face data and decent animation of face data is hard. Acquisition of human motion on set or on a soundstage is expensive and therefore is only done if absolutely necessary. It is hard to get good motion data from shooting on set, but this is required for addition of some post-production effects. The recent movie, *Avatar*, has set a new standard for digital humanoids and for animation of facial and motion-captured data. We note that there is still an enormous amount of work required in post-production and that, even now, the (pink-skinned) digital humans still do not look quite right in still imagery [Framestore].

g) **Breaking free from pixels**. A non-pixel format would be useful to break free from the problem of producing the same material at many different resolutions

---

[3] Lenny Lipton, "Digital stereoscopic cinema: the 21st century", *Proc. SPIE* **6803**, 2008, dx.doi.org/10.1117/12.759156

[4] Ronen Barzel, "Faking Dynamics of Ropes and Springs", *IEEE Computer Graphics & Applications* **17**(3), pp. 31–39, 1997, dx.doi.org/10.1109/38.586016

[5] F.C. Gee, W.N. Browne, K. Kawamura, "Uncanny valley revisited", IEEE International Workshop on Robot and Human Interactive Communication 2005 (ROMAN 2005), pp. 151–157, 2005, dx.doi.org/10.1109/ROMAN.2005.1513772

and needing to ensure that the original material is always shot at the highest resolution that you will need. Such a format would need to be able to handle all the processing that we currently do on sampled images. In the long term there would need to be input (cameras) and output devices that could handle the non-pixel format. For our own research plans, we would like to investigate further such a resolution-independent image format[6]. However, in order to solve the problem of generating a reliable movie format of this kind we will have to solve the motion vector analysis problem first (indeed a successful encoding gives us a solution). Fortunately it turns out that the single image version of a candidate format makes it equally easy to find motion vectors from any camera or object movement, so it may be possible to make progress both here and in the downstream side of motion vector analysis, e.g. recovery of 3-D data. [Smoke & Mirrors]

h) **Making non-photorealistic rendering useful**. Non-photorealistic rendering is not used much because the well known "shower door effect"[7] is difficult to avoid [Framestore, Cinesite].

i) **Novel user-interfaces**. We were somewhat surprised to find that there was no particular interest in the industry in new interfaces. In fact industry expressed concern that academics could consider interfaces to be a key issue. There have been some efforts in recent years to promote novel input devices (e.g. Monkey and the Dinosaur Input Device) but they have in practice turned out to be unusable for reasons which should have been exposed by a thorough HCI evaluation process. These devices essentially modelled simulacra of the objects to be manipulated but all lacked a feedback system which could have been used to drive the model to a given position for editing. This lack essentially made the device useless in practice but including it would have required costly precision engineering challenges which have not been fully resolved within the robotics industry (e.g. 'play' on extremities; safe, flexible and practical provision of considerable amounts of electrical power). The fact that these attempts at novel input devices have all failed is something which is not reflected in the survey literature, so could allow researchers to be misled [Framestore].

j) **Computer vision**. What we learned here came as something of a surprise too, but in the opposite way. Two years previously some of us had concluded that the film industry had little or no interest in computer vision technology. The accepted view was that computer vision was unable to provide algorithms which delivered the same quality of result as could be achieved laboriously by hand. In some difficult cases the state-of-the art in computer vision was well behind the needs of the industry. Computer vision techniques had only made inroads into problems that represent a tiny part of the overall technology, in particular camera positioning and object tracking. However we heard from several sources of a renewed interest in computer vision arising from the box-office successes of a number of 3D (stereoscopic) movies.

---

[6] John Patterson, Philip J. Willis, "Image Processing and Vectorisation" International patent application PCT/GB2007/002470, filed 5 July 2007, U.K. filing 3 July 2006.

[7] Amit Agrawal, "Non-photorealistic Rendering: Unleashing the Artist's Imagination [Graphically Speaking]", *IEEE Computer Graphics and Applications* 29(4):81–85 (July/August 2009)

One of the basic techniques in film effects, now widely used whenever digital intermediate is employed, is matte photography, usually known as 'travelling matte.' The process of combining images defined by travelling mattes is known as 'compositing.' Actors are filmed in front of blue or green[8] screens and their performances combined with CGI elements and with CGI or painted sets. If one wants to employ the same process in stereo imagery, then each version of the image elements has to be viewed or rendered correctly. This requires that correct depth of field, focal length, illumination, perspective and vanishing points have to be established for each member of a stereo pair and maintained throughout the sequence[9]. In order to determine these quantities consistently the compositors need to know about, and use, a lot more computer vision technology than hitherto. The problem remains that computer vision is no better at solving the film industry's problems than two years ago. For example, one of our contacts expressed his dearest wish as a reliable means of determining motion vectors for monocular movies. Motion vectors indicate the direction of motion of each element of the scene and would be extremely useful in effects work. If these could be extracted reliably, then one could carry out a comprehensive dense 3D analysis of a scene from motion vectors and camera position alone. The state-of-the art for 'easy' scenes is a reliability of 85%, at best, on translation alone (dolly shots). But even the state-of-the-art is hopeless on zooms and rotations, where 40% is a good result. Pans tend to be somewhere in between. These figures are simply not good enough for film effects.

2. Infrastructure issues

a) **Trans-coding media between digital formats**. There has been a proliferation of formats, which means that, for example, one advertisement can be required in 10 different formats. To compound this, different subsets of those 10 will be required for each country in which the advertisement is used. Further, a contract may be for up to 100 advertisements. The net result is that a lot of CPU time and staff time is spent in converting between video formats. A research aim is therefore to develop a video version of Adobe's Portable Document Format: a single file format that can be distributed and then converted at need at the player (see 1(g)). [Smoke & Mirrors]

b) **Transmission of large quantities of data including backup of large data stores**. A post-production or visual effects house will produce gigabytes of new data each day. One company reported that no vendor of off-site backup was able to cope with the quantity of new data that they produce. Two companies commented that, because of this, they maintain their backups on site, with the obvious security risk. There are also issues of transmitting large amounts of working data between a company's multiple sites (e.g. between the

[8] Other colours may be used for non human subjects.

[9] In practice it would be more usual to use stereo cameras (i.e. two eye cameras) precisely because such an approach finesses these difficulties. Then, in post-production, it is 'just' a matter of ensuring that these are kept together and synchronous. Systems such as Nuke™ or Quantel Stereo 3D™ already handle this form. Of course CGI elements will need the identification of these 3D parameters from the footage if combined with live action and this has to be done ahead of rendering into composable elements. It is at this point that the reliability of the, now unavoidable, computer vision techniques are being tested.

UK, the US, and China) and of transmitting large amounts of image data to customers. For example, one company uses 600 Mb links between its UK and US offices and moves image data around on Sohonet (i.e. locally[10]) at 1024 Mb. To give some idea of the scale of the problem, a movie frame will be several megapixels (maybe as much as 144 Mbytes), while an advertising poster can be rendered with up to 600 megapixels (1.8 Gbytes). One company has a 160 TB file store for handling its post-production work; another company mentioned data volumes of several hundred terabytes. [Smoke & Mirrors, Filmlight, Framestore]

c) **Keeping up with technology**. In parts of the industry, the basic algorithms have changed little in the past decade. The key problem faced in these parts of the industry is making best use of new technology to speed up processes and to keep ahead of the competition. For example, one company reported that only 10-20% of their code performed image processing, with the rest of the code being required for data management. [Filmlight]

d) **Archiving and cataloguing assets**. Archiving everything is problematic. If we do archive then cataloguing is important so that we know where to find things. For example, *The Tale of Desperaux* has 1700 effects shots, with 4 million assets, with variations on those assets producing 10 million identifiable objects (see Table 1). These take up several hundred terabytes. How do you archive something like this? How do you manage the archive? There are many subsidiary questions within this problem: for example, is it sufficient to store the original imagery and models along with a description of the process to get from those to the final shot? [Wavecrest, Framestore]

e) **Archiving footage in perpetuity**. There is a desire to archive the finished product forever. All physical media deteriorates, whether physical film, magnetic tape, or optical disc. Film has a life of around 40 years, though this varies considerably with storage conditions[11]. Some film has survived reasonably intact over 70 years[12]. LTO Ultrium (½" digital archive tape) has a predicted life of 15–30 years, albeit guaranteed by the manufacturers for only 7 years[13]. Can we develop mechanisms that robustly store digital footage for decades or centuries? If so, can we automatically migrate existing film archives to secure digital media? This is not a small problem: the British Film Institute has an archive of 150,000 movies: a thousand million feet of physical film, which constitutes just one third of their total archive of footage[14]. More footage is generated daily: the Internet Movie Database[15] reports 6,886 feature films released in 2009. One anecdote is of a feature film whose final digital version was archived on to tape in 2002. When the sequel was made in 2008, a flashback was required to the 2002 movie. The digital tapes were discovered to

[10] All the UK companies we spoke to are on Sohonet.

[11] James M. Riley, *IPI Storage Guide for Acetate Film*, Image Permanence Institute, 1993, acetguid.notlong.com

[12] British Film Institute Mitchell & Kenyon Collection, www.bfi.org.uk/ features/mk/ (accessed 26 February 2009)

[13] Sun Microsystems LTO Ultrium tape cartridge specifications, nle.ch/dl/LTO.pdf (accessed 26 February 2009)

[14] British Film Institute National Archive, www.bfi.org.uk/nftva/ (accessed 25 February 2009)

[15] www.imdb.org

be unreadable and the shot had to be re-made from the original film negatives, with all of the visual effects and post-production re-done from scratch. [Wavecrest, Cinesite]

f) **Healing the 2-D/3-D divide**. There are currently separate workflows for 2-D data (images) and 3-D data (modelling). It would be useful to join the workflows in some way, especially as stereoscopic movies become more popular. For example, one company's workflow has the 3-D lighting people passing rendered images to compositors for integration into the scanned film frames. The compositor then often has to adjust the image so that it sits well in the filmed frame, often to better approximate the imperfections of the film camera, or to adjust colours. Some shots require considerable adjustment, but the compositor only has the image to work with. Is it possible to increase the scope for these adjustments by integrating the 3-D and 2-D software? Can this be done by integrating a simple renderer into the composite graph? A simple example of where this is useful is in generating contact shadows, where a CG element needs to appear to rest on something in the live plate. This is awkward to do in image space, but it is a real help in speeding up integrating the elements, as it avoids the compositor needing to get the 3-D artist to produce additional renders. Packages like Nuke provide the hooks for this, but there is scope for more work in this area. [Foundry, Framestore]

g) **Improving digital capture**. There are currently no digital capture devices that can compete with film in quality of captured imagery. In particular CCD cameras cannot capture the dynamic range of film at maximum resolution (4K) [Cinesite]

h) **Improving digital image quantisation**. The standard storage format (DPX or Cineon) has been 10 bit logarithmic. The latest film stocks (since early 2006) have better quantisation resolution than this. Can the industry move to a new standard? If so what should it be? Logarithmic or linear? How many bits? Currently 32bit floating point is used as a default. [Cinesite]

i) **Managing intellectual property rights**. It was suggested to us that the rights model must change otherwise the industry will be strangled [Wavecrest]. An example of the degree to which the film industry (in this case, the US MPAA—Motion Picture Association of America) currently protects its rights extends to a scheme of certification available to sites which handle film for the studios. This requires that film assets are only accessible by authorized personnel and the access points be monitored by security cameras. The MPAA also requires that the digital files be isolated on localised networks where studio owned images are moved electronically. This extends to monitoring such devices as film cleaners, viewing machines and rewinding benches. What the MPAA is doing is offering what is in effect a security certification[16] whose standards are published[17] and their studios/producer members may choose to work only at places that have been certified, although this is in no way an imposed requirement. The MPAA process is a means by

[16] In practice for our hosts this consisted of a very thorough audit in 2008 after which MPAA issued a very detailed "Site Survey Report" which is also given to their members. They didn't choose to be certified as such – they were notified of an audit and they like to oblige. In essence this amounts to the same thing as certification for our host, as the 2008 report was favourable. Our hosts also have a FACT security certificate.

[17] See http://www.mpaa.org/piracy.asp for a statement of the policy behind the MPAA's standards.

Item	Statistic
Number of shots	1713
Number of Locations	63
Number of (hero) characters	53
Number of variants in crowds	263
Number of props	1080
Number of 3D models	6098
Crew Size (peak)	280
Render Farm (CPUs)	4500
On-line data	150 Tbytes
Number of published versions of assets	4,031,382
Dependencies between assets	20,375,436
Metadata (number of objects)	29,797,895
Metadata (number of attributes associated with objects)	397, 714,992

Table 1: Statistics from the making of *The Tale of Despereaux* [Framestore]

which companies can demonstrate they take the utmost care of rights handling. Many facilities in the UK are not so certified and still work on mainline films. However the big four effects companies have chosen to be certified so that they are automatically eligible for all work. The MPAA procedure is a rigorous and extends beyond internal security.

In practice, some studios go far beyond the MPAA specification. Some productions, for example, refuse to allow any of their material to be transferred onto portable disc: and so everything had to be synchronised round the effects houses using Sohonet [Wavecrest, Cinesite].

j) **Data volumes**. As listed above (a,b,d,e), we were advised about a number of problems the industry is facing due to the sheer volume of data they have to deal with on a daily basis. The scale of the problem is daunting. The culprit is digital intermediate, where every image in a film is stored as a file in a database. The problem is exacerbated for 3-D stereoscopy because the storage doubles for every element. The ease with which these images can be manipulated in real time allows directors to try out lots of ideas for, say, colourisation and the apparent ease with which they can be stored encourages these same directors to keep every version until the final edit, maybe as many as 50 versions. Once upon a time, the master negative was on film and only edit prints (which were struck once from the master and then from each other) were available for this kind of treatment and so experimentation was minimal. Today, the workflow model is moving towards far more asset-heavy CGI work, not only for CG feature animation but also for visual effects for feature films, all facilitated by the ease of introduction and manipulation of digital assets in digital intermediate format. One of our hosts [Framestore] identified CG-heavy effects work as being the more obvious source of pressure than those directly attributable to digital intermediate itself. They had just completed a CG feature animation (*The Tale of Despereaux*) and had publicly released some statistics from their work here which makes the point (see Table 1). There is also a lot of pressure in the industry to use digital intermediate because of the insurance premiums for the master film negative. If the master negative is lost or destroyed at a late stage in filming then the entire

cost of making the film has to be recovered from the insurer. Insurance premiums of two million US dollars are not unknown for the master negative. By contrast digital intermediate allows for perfect digital copies and the entire database can be duplicated, allowing equivalent premiums of $50,000. On the other hand, one needs to be able to capture a digital image to the same quality as is obtainable from a film camera and this was also an issue with two of our companies (see 2(g)). At present this can only be achieved using film and then scanning in the film a frame at a time at about five seconds a frame.

3. People and process issues

a) **Managing artists**. A decade ago, the artists in the industry were generally aware of the underlying technology and of the entire pipeline of getting from concept to the finished film. Today, the artists are often less technically knowledgeable. Because they specialise more in particular roles they frequently understand less of the technical underpinnings of the tools they use and they therefore can either fail to use the full power of the technology or fail to understand the implications of their actions for the later stages of the pipeline. [Framestore]

b) **Managing client expectations**. Much visual effects work is time-consuming and labour-intensive. Many effects are generated using one-off solutions that are hacked together to get the result wanted by the director. Despite these difficulties, the companies find that their clients have little appreciation of which effects are straightforward to produce and which are extraordinarily expensive. There is a common believe that, if they have seen an effect in some other movie, then it must be straightforward to produce. Indeed, "special effects" are now an ordinary part of the production pipeline and are not "special" at all. Some movies have over a thousand effects shots and even non-effects movies employ a lot of digital post-production. *Mamma Mia*, a live-action movie, had over 900 shots that required post-production, for example, changing the sky colour and moving or removing background elements. There is also the problem that clients often cannot specify exactly what they need, requiring multiple re-rendering or re-formatting. [Framestore, Smoke & Mirrors, Cinesite]

c) **Managing a large workforce**. The industry once consisted of small companies within each of which everyone knew everyone else. Over the last decade, several of the companies have become too large to work in this way. How do we manage this creative, collaborative process when people in different parts of the chain do not know each other and have only a basic understanding of each other's roles? [Framestore]

d) **Managing workflow**. How do we break out of the production line method of producing effects and post-production? How do we provide effective feedback loops between the different links in the production chain? [Framestore, Foundry, Cinesite]. Creativity is driven by iteration—going over the same sequence again and again trying out different ideas—the industry needs tools to manage iterative procedures [Cinesite].

Conclusion

We believe that we have successfully plumbed the major research issues and concerns of the post-production and visual effects industry in London. We have subsequently passed our report to two other major houses, who have recognised many of the problems as being ones that they also face.

Computer graphics and image processing researchers, like us, are best placed to tackle the development of new technologies in (1). These are also the problems best suited to university timescales. We are working with some of the companies to research certain of these. Our colleagues in networking, information retrieval, databases, and engineering are best placed to tackle research issues in infrastructure (2), particularly how to handle backup and archive of large datasets. The managerial issues (3) demonstrate that some of the biggest problems facing the industry have little to do with technology and everything to do with people.

Final comment: the importance of the industry

The film industry has long suffered from an attitude that it is somehow 'insignificant', that its problems are not worthy to be a subject of academic research, or even that its outstanding problems were solved long ago. This tends to be an academic attitude. The fact that someone has published a solution to a simplified version of a particular problem in, say, vision, does not mean that solution will scale up to the requirements of the film industry. In fact, it usually fails to scale, resulting in distrust in the peddlers of technology on the industry side also. We were gratified to be kindly received by everyone in the industry we asked to see, indeed by rather more people than we anticipated so we had a busy couple of days running round Soho. We thank everyone we met for receiving us, discussing the issues with us so comprehensively, and allowing us to spend so much time with them. We usually asked for 60 minutes but we always got more than that.

We were pressed by one of our hosts to point out that in the last quarter of 2008 three British films, *Quantum of Solace*, *Dark Knight*, and *Mamma Mia* between them made over £2 billion at the box office. They will likely generate revenues in the vertical market (the music track, the DVD, the game, the T-shirt, the merchandising, spin-off books) of 5–10 times that amount. The films were British in the sense that they were filmed and post-produced here but the bulk of the revenue is repatriated to the film's financial backers in the USA [Cinesite].

This is not an insignificant industry, and adequately facing its challenges is important for the UK economy.