

INTRODUCTION

The principle at the heart of moving pictures is the ability to project still images at the speed of 24 frames per second, thus creating the illusion of continuous movement. Animation takes advantage of the fact that the filming of such images need not be continuous, and that each frame in a sequence can be photographed individually with any amount of time between exposures. By manipulating objects incrementally from frame to frame during this time, animators are able to give inanimate objects the appearance of movement and life when the film is ultimately projected at the normal speed.

There are two broad categories of animation. 2-D animation mainly involves the photography of flat artwork such as drawings or paintings to produce what we might call a cartoon. Such methods are used to produce short films featuring popular characters such as *The Simpsons* or, in their most complex form, animated feature films such as *Pinocchio* (1940) or *Beauty and the Beast* (1991). Less well known is the practice of animating 2-D artwork to create visual effects for live-action films. Elements such as sparks, lightning and laser bolts can all be drawn by hand and photographed one frame at a time before being added to live-action images. These traditionally painstaking and laborious processes have been greatly affected by the emergence of the computer, but the principle of animating 2-D images remains essentially unchanged.

3-D animation processes, on the other hand, involve manipulating the physical position of dimensional objects such as puppets and environments from frame to frame. Such techniques are best known for creating the performance of fantastic creatures for films such as *King Kong* (1933) or *The 7th Voyage of Sinbad* (1958). 3-D animation has also been greatly affected by the digital revolution and many stunningly realistic images, from dinosaurs to spaceships and even clouds of dust, are now created and animated entirely within the computer.

Animation, in all its varied forms, is perhaps the most important and widely used of all special effects techniques and has enabled the production of some of the most memorable and breathtaking moments in the history of the movies.



2-D ANIMATION

It was a decade after the emergence of moving pictures that 2-D animation, the ancestor of the modern cartoon, became a recognizable mode of production. In 1906 J. Stuart Blackton (1875–1941) produced a film called *Humorous Phases of Funny Faces*, in which an artist's hand sketched faces in white chalk on a black background. Once drawn, the caricatures assumed a life of their own and performed exaggerated facial expressions. The film was made by a process of exposing a frame or two of the drawn face, then redrawing the face in its new position and exposing it for another frame or two.

In 1908 the French caricaturist and comic-strip artist Émile Cohl (1857–1938) made *Fantasmagorie*, a two-minute film of hand-drawn, moving stick figures. Between 1908 and 1910, Cohl contributed animation to over 75 films for the Gaumont company, and invented many of the tools that became standards of the trade, including the vertically mounted camera animation stand and charts for plotting character and camera movement. Cohl's surreal animations, in which fantastic drawings metamorphosed into one another with no apparent logic, were highly successful, and in 1912 Cohl moved to the United States to produce a series of animations based on a newspaper comic strip called *The Newlyweds and Their Baby*.

A number of American newspaper cartoonists were encouraged by the success of Cohl's films and began to make animated adaptations of their own work. Among them was Winsor McCay (1871–1934), who drew the popular *Little Nemo* comic strip for the *New York Herald*. In 1911 McCay drew and filmed 4,000 images to produce a short film called *Little Nemo*. Before the animation itself, McCay appeared on screen to demonstrate the way in which thousands of individual drawings were photographed to create the impression of movement. McCay actually painted directly onto the black-and-white celluloid images, to produce the first colour animated film.

After a year's work, McCay then produced *Gertie the Dinosaur* (1914) – not, as is often claimed, the world's first cartoon, but certainly the first to feature a character with a personality. Gertie appeared to move with a feeling of ponderous weight, was shy but liked to show off, and could even shed tears. The film was presented by McCay himself, who would stand to one side of the screen giving instructions to which Gertie appeared to respond. To create the film, McCay drew every detail of every frame on thousands of sheets of paper. The paper was semi-transparent, enabling McCay to trace the previous image and make the slight differences that would produce movement when the pictures were filmed and shown in rapid succession.

In 1914 the cartoonist and producer John Bray (1879–1978) pioneered a method of painting the background scenery of a shot on a clear cellulose acetate sheet, leaving the area where the moving characters were to appear unpainted. The acetate sheet was then placed on top of the characters – still



drawn individually on paper – and filmed. The system made production much quicker since only the moving characters and not the static background scenery had to be redrawn for each shot. Bray's system was improved later in the same year by Earl Hurd (1880–1940), who more logically reversed the process, using a single elaborate background painting over which cellulose sheets or 'cels' containing changing images of the characters were laid. The method of producing characters on see-through cels has formed the basis of 2-D animation ever since.

Realizing the potential of animated cartoons, the Edison company financed what is thought to be the first cartoon production company, Barré Studio, under the supervision of Raoul Barré (1874–1932) and William Nolan (1894–1954), who together devised several techniques that became industry standards. Around 1914, Nolan and Barré added a row of pegs to drawing boards and animation stands, on which cels and backgrounds with perforations could be laid to ensure perfect registration during drawing and photography. They also introduced an assembly-line production regime in which separate teams of people drew animation, painted it, photographed it and so on. Even the process of drawing animation itself was streamlined, with lead animators sketching characters' most extreme poses, while lower-paid artists drew the actions in between – a process that became known as 'tweening'.

By the 20s, the animated cartoon was an established part of every movie theatre programme, and animated characters such as Koko the Clown and the incredibly popular Felix the Cat were as widely recognized as their live-action co-stars. In 1923 a young and ambitious animator moved to Hollywood from Kansas City, where he had been producing humorous animated advertisements called 'Laugh-O-Grams' for a local theatre chain. Within five years he was producing the most popular and sophisticated cartoons in the world, and his name soon became synonymous with animation.

Walter Elias Disney (1901–66) and his partner Ub Iwerks (1901–71) produced, among other projects, a series of *Alice* comedies (1923–7), technically advanced cartoons that combined a live-action character called Alice with a cartoon world and cartoon characters. Disney and Iwerks made over 50 *Alice* films before creating a new character called Oswald the Lucky Rabbit in 1927. The cartoons starring Oswald were a success, so much so that their distributor, Charles Mintz, poached most of Disney's staff and set up his own studio to produce the films without Disney's input. Disney vowed that he would never again make films for anyone else. Looking for a new character to revive his fortunes, Disney shortened Oswald's ears, gave him a long tail and turned him into a mouse called Mickey.

Disney produced two Mickey Mouse cartoons in 1928, but, unable to find a national distributor, decided to make a third with sound. *Steamboat Willie* (1928) was not the first sound cartoon; Oswald the Lucky Rabbit had already been set to music, as had the work of Max and Dave Fleischer (174>) and the Terrytoon cartoons made by Paul Terry (1887–1971). However, *Steamboat Willie* was the first cartoon to use synchronized music and sound effects, including a whistling mouse and cats that miaowed when their tails were pulled. Audiences and critics loved the seven-minute cartoon and demanded more. Disney never looked back. In 1932 Disney won an Oscar for *Flowers and Trees*, the first cartoon made in three-strip Technicolor (<56), and went on to collect every Academy Award for animation over the next 11 years.

While other American animators were happy to produce cartoons that portrayed a slapstick, almost surreal version of the world, Disney developed a unique form of stylized realism combined with a close attention to narrative and characterization. By the mid-30s, Disney felt that the art of animation had developed sufficiently to support the production of a feature-length animated film. In 1934, against the advice of almost everybody, Disney began production of *Snow White and the Seven Dwarfs* (1937).

PRECEDING PAGES: **The mighty King Kong** (1933), one of cinema's first great animated characters, and still one of the best loved.

ABOVE LEFT: Simple hand-drawn stick figures were animated in Émile Cohl's early cartoon *Fantasmagorie* (1908).

LEFT: **Gertie the Dinosaur** was one of cinema's first cartoon stars. Winsor McCay drew every frame by hand, even creating a cartoon version of himself to interact with the characterful dinosaur.

RIGHT: Enjoying the new medium of sound, Mickey Mouse uses a cow's teeth as a xylophone in Walt Disney's pioneering cartoon *Steamboat Willie* (1928).






FIGURE 1 THE MULTIPLANE CAMERA



composite image
as seen by camera



individual elements
arranged in layers

lighting for
each layer

Disney knew that the limitations of normal 2-D cartoon animation would constrain the aesthetics of the feature-length film that he had in mind. Disney's main concern was to find a way of making his cartoons look less flat and two-dimensional. In the real world, a camera that pans over any scene will make objects in the distance appear to move very slowly, while those in the foreground appear to move much more quickly. This phenomenon is a vital cue to the depth of a scene and is known as 'parallax shift'. Since the painted cels used in ordinary animation are placed directly on top of one another with no distance between them, there is no shift in parallax when the camera moves over them. Furthermore, real scenes of objects that are some distance apart are affected by depth of field (<115), which makes foreground and background elements appear more or less in focus. Sandwiched acetate cels have no distance between them, so every layer remains in constant focus.

To overcome the lack of depth in ordinary animation, the Disney technical department built the first multiplane camera (fig. 1) – an award-winning and revolutionary piece of animation equipment. The camera was a towering device in which several animation cels were layered vertically with some distance between them. The camera, which looked down onto the stack of cels, was then sharply focused on a single layer so that the elements in front and behind were in soft focus. The cels could also be moved horizontally so that particular elements, such as the branches of overhanging trees, could be moved out of the way to give the impression that the camera was moving through a scene. Lengthy tracking shots were achieved by moving long sheets

of painted acetate or glass past the stationary camera. The cels in the foreground were moved much faster than those in the background to give a convincing sense of depth. Each layer of animation was individually illuminated to enable the creation of interesting lighting effects.

Using the multiplane camera was a complicated process. The movements of each of the separate layers of animation were painstakingly calculated and plotted before photography began. Complex sequences involving the movement of multiple layers, sometimes with dozens of characters, could involve using a team of animation photographers to service the requirements of each layer of imagery. After days of work, the success of a shot would only be confirmed when the developed film was returned from the laboratory.

Before being used for *Snow White*, the multiplane camera was given a trial run for the short film *The Old Mill* (1937), for which it created richly atmospheric images with a sense of depth that had never been seen in 2-D animation before. The incredibly realistic feeling of depth in *Snow White* was partly responsible for the film's massive commercial success when it was released in February 1938. Walt Disney had proved to his critics that the public would pay to see feature-length animated films, and his studio embarked on the production of an ambitious and profitable series of animated features, many of which have become classics that are enjoyed to this day. Though Disney's storytelling became increasingly sophisticated, the place of the multiplane camera at the heart of the animation production process remained unchanged until the arrival of computer-aided techniques in the 80s.

SPECIAL EFFECTS ANIMATION

Although creating the illusion of life with a few strokes of a pencil could be considered a special effect in itself, traditional 2-D animators distinguish between character animation and special effects animation. Anything that moves that isn't a character is usually animated by the effects department. This can include rain, snow, water, falling leaves, fire, smoke, shadows – anything that brings the scene to life.

As with so many aspects of animation, Walt Disney was an early innovator in the field of special effects. Character animators were traditionally responsible for drawing their own special effects, and as a result, such elements often remained crude and ineffective compared to the advances that were made in stylistic character animation. Disney was aware that his first feature film, *Snow White and the Seven Dwarfs* (1937), would rely just as much on the credibility of its environment as it would on its characters. To create the world he envisaged, Disney established the first special effects animation department and staffed it with artists whose talents lay in the portrayal of the natural world. These animators spent months studying the way in which different types of rain fell, how rivers flowed and flames danced. With nothing to animate but special effects, the skills of the effects department staff developed so greatly that, as *Snow White* neared completion, the effects animation created at the beginning of production had to be redesigned to match the sophistication of the later work.

One of the major innovations made by the Disney effects department was a new method of creating complex shadows. Traditionally, characters looked rather flat because there were no shadows on their bodies. The shadows that characters cast on the ground and the objects around them, which are so important in tying a character to its environment, were usually limited to amorphous dark patches that randomly followed the characters around. For *Snow White*, the Disney effects artists completed the character animation and painting, then placed another cel on top of each character on which they drew the outline of a shadow area. The shadow area was then painted solid black. During photography, the shadow cel was laid on top of the character cel and the two were photographed for around 70 per cent of the time required for a good exposure. The shadow element was then removed and the character painting was photographed on its own for the remaining 30 per cent of exposure time. The result was an image that contained a character with a transparent shadow area – the character detail

could actually be seen beneath the shadow. By varying the proportion of exposure time given to the shadow element, shadows could be made darker or lighter depending on the mood and lighting of the scene.


'Shadows continue to be a large part of an effects animator's job,' says Jon Brook, an animation effects supervisor who has worked on modern animated features such as *Who Framed Roger Rabbit* (1988) and *An American Tail: Fievel Goes West* (1991). 'You might think that creating a shadow down the side of a character is simply a case of drawing an area of darkness over the body, but in fact, shadow style varies greatly from scene to scene and film to film, depending on the overall production design. Shadows can be hard-edged and bold, or soft-edged and subtle so that they wrap themselves gently around the body. Shadows also have to be animated according to the content of the scene; as a character moves past a light source, the direction and quality of a shadow will change. If there's a fire in the scene, a shadow will shudder and move in response to the mood of the fire.'

Perhaps the greatest challenge for an effects animator is water. Water is transparent, elastic, constantly moving, sometimes heavy, sometimes light and frothy. Re-creating water with a few hand-drawn pencil lines is an extraordinary skill. A wave is not simply a crest that can be drawn with a single line – a wave has shadows, currents and smaller waves that grow and die within it as it moves. 'A simple little splash has so much thought put into it,' says Brook. 'It has to grow and die in probably less than a second, yet in that 24 frames we must draw a series of droplets that separate out and rise up to a crest before falling back down convincingly, all within the specific style of the production. We don't simply draw a standard splash when one is needed; we spend a great deal of time designing the feeling of the splashes in any production so that they fit in with the overall art direction – they might be hard, spiky splashes or delicate, soft ones.'

An effects animator is often called upon to draw objects as they move and change position. A classic example is leaves blowing in the wind. Each leaf must change its shape, colour and shadow as it is twisted and carried by the breeze. Some of the most extraordinary examples of this artistic skill can be seen in Disney's *Fantasia* (1940). For the Nutcracker Suite sequence, animator Cy Young created a beautiful moment when a delicate white blossom drifts down to land on the surface of a pond before being reborn as a graceful ballerina swirling up into the sky. This moment is one of hundreds of sublime animated effects in the film and, despite being over half a century old, *Fantasia* remains perhaps the finest example of the art that has ever been produced.

FAR LEFT: Invented in 1937, Disney's multiplane camera brought a new sense of depth and realism to 2-D cartoon animation. The vast monolith took a large crew to operate and required each scene to be plotted in minute detail – but the results were spectacular.

RIGHT: Examples of irregular movement, particularly natural phenomena such as splashing water, are painstakingly designed and drawn by special effects animators. These production designs were created by animation effects supervisor Jon Brooks.





ROTOSCOPING

The use of effects animation is not confined to films featuring cartoon characters. Many live-action feature films make extensive use of hand-drawn 2-D animated special effects.


In 1917 the animator Max Fleischer (1883–1972) patented the rotoscope, a device that projected pre-filmed footage of a performing actor onto a sheet of glass. The movements of the actor were then traced onto sheets of paper one frame at a time, and these were used as templates to draw cartoon characters that, when animated, had incredibly lifelike movements. Fleischer and his brother Dave (1894–1979) used the rotoscope to produce their popular *Koko the Clown* cartoons of the 20s and the *Betty Boop* and *Popeye* cartoons of the 30s. The rotoscope did not revolutionize the production of cartoons in the way that Fleischer had hoped, though it has been used to produce character motion for a number of films including Ralph Bakshi's animated version of *The Lord of the Rings* (1978). Most character animators, however, have preferred to use a stylized hand-drawn form of human movement.

Though rarely used in the production of animated cartoons, rotoscoping has found other important uses in the field of film special effects. By projecting pre-filmed images onto a flat surface, they can be traced to

produce hand-drawn travelling mattes (<67). The process can also be used to hand-draw a variety of 2-D animated elements, such as lightning, lasers, gun blasts and shadows.

Steve Begg is one of Britain's leading special effects artists and is recognized for his skill in hand-drawing animated special effects elements. 'One of the most commonly required animation effects is lightning and electrical charges,' says Begg. 'To produce such effects optically, we first load a registration print [the developed print of the sequence that needs an animation effect added to it] into the rotoscope projector. Each frame of the sequence is then projected down onto a piece of paper. Using the projected image for reference, we hand-draw a rough version of the actual animation itself, or trace reference points from the scene that we can use as a guide for our animation later. The artist then sketches the animation in pencil before going over it with black ink or paint.

'There are a number of ways of getting lightning to look real,' says Begg. 'Some people like to draw the tip of the bolt of lightning as it emerges from a cloud and then extend it down towards the ground over three or four frames. For natural lightning, I think it looks best if it snaps down to the ground instantly. If it's some sort of fantasy film and you want the lightning to appear to have a life of its own, you can draw it so that it twists



FAR LEFT: The mighty dragon Vermithrax is struck by hand-drawn lightning effects created by ILM for *Dragonslayer* (1981).

LEFT: Effects supervisor Steve Begg photographs flat artwork using the rotoscope at The Magic Camera Company.

BELOW: Actors fought with wooden sticks during the making of *The Empire Strikes Back* (1980). The glowing light sabres were later hand-drawn and painted before being added to the footage during compositing.

and curls in an organic-looking way before recoiling back into itself. It's important to try to give it that life.'

Once the animation of lightning or lasers has been drawn and painted, it is photographed. Each sheet of animation is laid back on the rotoscope stand and photographed sequentially onto high-contrast black-and-white film. When developed, the result is a series of black negatives with white (transparent) animation on them; in other words, the black-painted lightning becomes a clear area on the negative. This film is then loaded into an optical printer (<70) and combined with the film of the original scene. 'We will often print animation in several passes', explains Begg. 'Drawn animation is usually quite sharp-edged. This is fine for sharp electrical sparks and lighting, but lasers and other effects often look better if they have a soft glow. To achieve this, we first print the animation in what we call a "core pass". This gives a bright, sharp centre to the effect. Then we rewind the animation and print it a second time using a diffusion filter that blurs the image. This produces a soft glow around an intense centre. By placing a coloured filter in front of the animation during printing, the lightning or laser bolts can be made any colour that you want.'

As well as creating areas of light, effects animation can be used to create areas of darkness. When real aircraft fly over a real landscape, they cast shadows beneath them as they travel. When model aircraft are composited into real or model landscapes, the two elements have no physical relationship and no corresponding shadows. By using the rotoscope to plot the path of a flying object, shadows can be drawn by hand that trace the shape of the aircraft on the contours of the ground below. By adding the shadow of one object onto another, the separate elements of a scene are effectively 'tied' together, as if filmed at the same time. Fine examples of this method can be seen in *The Empire Strikes Back* (1980). When the *Millennium Falcon* is pursued across the surface of an asteroid by Imperial TIE fighters, the spaceship's shadow can be seen rippling over the rocky surface below.

Rotoscoped effects animation can also be used to create more down-to-earth effects. 'Over the years there has been one simple effect that we have been asked to do probably more than any other', explains Gene Warren, head of Fantasy II Film Effects. 'Lots of films contain shooting, and getting a gun to look as if it's firing can be a hit-or-miss thing. Firstly, the person using the gun has to remember to pull the trigger – in all the excitement of an action scene, they sometimes forget to do this. Secondly, even if the gun does go off, the flash of the gun lasts for just a fraction of a second and it is quite possible for it to occur between frames when the shutter of the camera is closed. We often end up having to rotoscope the muzzle blast on guns into a scene; we've done this hundreds of times over the years.'

'We do muzzle blasts just like normal rotoscoped animation', explains Warren. 'The shot is projected down onto paper and the position of the necessary flash is traced and then shaded in pencil. Pencil shading is great because it gives you a random, grainy look that is perfect for muzzle blast and it only takes about five minutes to draw. The drawing is then photographed and printed into the shot in the optical printer. The last little touch is to place a filter in front of the image during optical printing and put a dab of petroleum jelly onto the filter. With a fine brush, we just draw the grease out along the path of the bullet – this gives an almost imperceptible blur as if a bullet really is coming out of the gun. One of the biggest jobs we did this for was *Point Break* [1991]. We did nearly all the gun blasts in the bank robbery scenes. A few of those muzzle blasts are real and lots of them are fake, but it's totally impossible to tell which is which – I'd have to go and check our records to find out which ones we did because even I can't see the difference!'

Today rotoscoping and effects animation is achieved digitally with artwork being either hand-drawn and scanned into the computer or painted directly in the computer using specialist digital paint software (<108).



WHO FRAMED ROGER RABBIT

Film-makers have always been fascinated by the potential of combining live-action performers with animated characters. Max Fleischer's (<174) *Out of the Inkwell* series of cartoons (1919–28) featured animated characters that left their artificial environment for a jaunt in the live-action world. Walt Disney (<171) reversed the conceit in his *Alice* comedies (1923–7), in which a live-action actor was placed within an animated environment populated by cartoon co-stars. Later feature films such as *Anchors Aweigh* (1945) and *Pete's Dragon* (1977) combined human and cartoon characters, but the blend was never particularly subtle.

When director Robert Zemeckis (<43) and executive producer Steven Spielberg (<39) decided to make *Who Framed Roger Rabbit* (1988) they knew that modern audiences would only watch the film if human characters and 'Toons' coexisted with a level of realistic interaction that had never been achieved before. After two years in production, the film that was finally released blurred the lines between animation and live action until they were almost unrecognizable. Audiences were thrilled by a world that included such inconceivable marvels as a cartoon car with a human passenger being chased by a real car driven by a gang of cartoon weasels.

The first stage of production involved filming the live-action scenes into which animation would eventually be inserted. Because the film had to go through numerous optical processes before completion, the entire movie was shot in the large VistaVision format (<55), which helped to retain good picture quality. It was the first film shot entirely in the format since the 50s.

Star Bob Hoskins endured months of reacting to non-existent characters. To help Hoskins, voice artist Charles Fleischer dressed in a Roger Rabbit costume and read Roger's lines from off-camera. Hoskins was also surrounded by dozens of technicians operating a host of remote-controlled props that had been cleverly created by physical effects supervisor George Gibbs (330>) to indicate the presence of an unseen co-star. Coats bulged with the contortions of an invisible rabbit, chair seats depressed and puffed out dust as imaginary characters sat down on them, and guns floated around in the hands of a transparent pack of weasels – Zemeckis referred to the proceedings as 'the most elaborate invisible man film ever made'.

Once live-action photography was completed, the film was edited into a semblance of its ultimate form. Scenes awaiting Toon performers were then enlarged into a series of photographs. These blow-ups were used as a guide by the 375 animators working under the supervision of master animator Richard Williams at his London studio. Williams produced rough pencil-line performances and combined them with the live-action backgrounds for the director's approval. Animation was particularly taxing because Zemeckis had filmed many scenes with a moving camera. To create cartoon characters that fitted into these scenes animators had to draw Toons whose perspective altered as if filmed by the same camera.

Once approved, animation was redrawn, inked and painted. To help bring a sense of 3-D life to the 2-D characters, a new method of producing shadows was devised. Since the 30s shadows had been painted in black and double-exposed into a shot (<173), which gave all the shadows the same grey, sometimes muddy tone. To help the Toons blend more realistically with their surroundings, all shadow detail was painted in coloured tones; Jessica Rabbit's red dress was fringed with dark burgundy shadows and scarlet highlights, for example. Additional layers of animated effects – sparkling sequins in the case of Jessica's dress – were also created. Extensive rotoscoping (<174) was used to produce hand-drawn travelling mattes (<67) so that Toons, humans and real-life props could be composited together as if actually interacting.

The 82,000 frames of hand-drawn animation were then optically composited at Industrial Light and Magic to produce over 1,000 shots – equivalent, according to visual effects supervisor Ken Ralston (299>), to creating the effects for all three *Star Wars* films. Upon its release, *Who Framed Roger Rabbit* was a huge hit, earning millions at the box office and an Academy Award for its innovative and painstakingly produced special effects.



SLIT SCAN

For the tunnel of light known as the 'Stargate sequence' in *2001: A Space Odyssey* (1968), director Stanley Kubrick simply told his effects crew that the camera should appear to 'go through something'. Several options were considered, but it was a proposal from Douglas Trumbull that caught Kubrick's imagination.


Experiments in abstract streak photography had been carried out as early as the 40s by brothers James and John Whitney. In the 50s John Whitney (<146) acquired an army-surplus analogue computer and rebuilt it to photograph patterns of light for an abstract film called *Catalog 61* (1961).

In the system developed by Trumbull with Con Pederson, Bob Abel and Colin Cantwell, an animation camera (fig. 2 (a)) was placed on a 4.6 m (15 ft) track, on which it could be moved slowly and smoothly backwards and forwards using worm gearing (b). Two large sheets of glass were placed at the far end of the track, each sheet about 1.5 m (5 ft) high and 3 m (10 ft) wide. These sheets could be moved vertically and horizontally. The sheet of glass farthest from the camera held transparent, backlit artwork (c). The sheet of glass nearest the camera was masked in black material, apart from a small slit (d), from which the system derived its name. This slit allowed only a small area of the backlit artwork to be visible to the camera at any one time.

To produce the streaking effect seen in the film, the camera began with its lens close to the slit so that the artwork behind it reached the edges of the frame. The shutter on the camera was then opened to expose a single frame of film to the image. Still exposing the same frame of film, the camera was moved away from the slit until it reached the end of the track. As the camera moved backwards, gears ensured that its lens kept the shrinking artwork in focus. At the same time, the artwork itself was moved horizontally or vertically across the slit to provide distortion. When the camera reached the point farthest from the artwork, the shutter was closed after having exposed, on one frame of film, a distorted streak of coloured light as it travelled away from the camera towards the middle of the screen (e). The process can be compared to taking photographs of roads at night, where long exposures register passing vehicles simply as the blurred streaks of headlights.

Next, the film was rewound and the process repeated with the blur of light travelling from the opposite side of the frame towards the centre. The process was carefully programmed so that for each subsequent frame filmed, the movements of the camera and the artwork were repeated with minute, sequential differences. The result was the effect of travelling through the light tunnel.


No hard-and-fast rules were followed when producing the sequence. Each filming session was set up with different variations of artwork, movement and exposure, and after many hours of laborious work, the result might be, according to Kubrick, 'like carpet going by' or one of the most dazzling, kaleidoscopic light shows ever to have been captured on film. After slit scan's spectacular debut in *2001*, it was used to create streaking title effects for films such as *Superman* (1978) and the wormhole sequence in *Star Trek* (1979).



ABOVE: The ultimate trip, courtesy of Stanley Kubrick's *2001: A Space Odyssey* (1968). The psychedelic light show of the Stargate sequence was created using slit-scan photography.

RIGHT: The use of tens of thousands of digitally animated 3-D characters gave a sense of massive scale to the Exodus sequence in DreamWorks Animation's *The Prince of Egypt* (1998).

FIGURE 2 SLIT SCAN



2-D COMPUTER ANIMATION

Everything about traditional hand-drawn animation is laborious and painstakingly slow, so when computers began to affect other areas of movie production, animation studios wasted no time in finding ways to streamline their own production methods.

In 1986 Disney, in association with Pixar, began to develop CAPS (computer animation production system). CAPS was designed to perform many of the tasks of the traditional animator within the computer. First, the black-and-white artwork was drawn by hand in the traditional manner before being scanned into a computer. Each drawing was then coloured on screen using a digital paint system. The rows of artists who had once worked amid pots of freshly mixed paint were replaced by a few computer monitors. While one artist could hand-paint perhaps 20 cels a day, a single computer operator could now complete around 200. As well as painting images, CAPS could also be used to generate certain special effects and to assemble multi-layer scenes without the use of a multiplane camera (<172). The first feature film entirely painted and assembled using CAPS was *The Rescuers Down Under* (1990).

2-D animators also wanted to take advantage of the computer's ability to create 3-D images. While traditional animators could draw scenes that replicated a complicated camera move, the process was highly skilled, time-consuming, and not always successful. In the computer, however, objects created in 3-D could have any movement or change in camera angle applied to them at the touch of a few buttons.

In the 80s several systems were developed to turn 3-D computer-generated objects into 2-D animation. These systems allowed relatively simple geometric shapes to be modelled, animated and 'filmed' in the computer. The resulting sequences could be printed on paper as a series of 2-D frames in which objects had black lines around their edges. These frames could then be individually painted and combined with traditional backgrounds and character animation. This technique was first used for some scenes in Disney's ill-received *The Black Cauldron* (1985). It was used more effectively for the climactic scenes of *The Great Mouse Detective* (1986) when it produced dynamic shots of Big Ben's ticking clock mechanisms.

Computers are now an essential part of the production process of 2-D animated feature films. At the forefront of modern animated feature production is DreamWorks Animation, whose ambitious first movie *The Prince of Egypt* (1998) was a landmark in feature animation. 'The Prince of Egypt was an epic story told on an epic scale,' comments 3-D layout artist Harald Kraut. 'We felt that traditional 2-D multiplane camera work – or at least today's computerized equivalent – wasn't up to the scope of the story. We didn't want all of the wonderful Egyptian architecture to look like pieces of flat artwork sliding past each other; we wanted to create a real sense of environment.' To help create this environment, DreamWorks developed a new animation system called 'Exposure'.

The Exposure system is based on the principles of normal 3-D animation production, in which a 3-D environment that has been built, textured and lit in the computer has a virtual camera (233>) placed within it. The camera is then animated to move around and 'film' objects from any angle. Using Exposure,





DreamWorks artists construct digital 3-D buildings and objects in the usual way, but rather than applying computer-generated textures (<164), the exterior of the objects is left flat and without details. 'We position the virtual camera to look at each object from an angle that shows the largest amount of surface area on that model,' explains Kraut. 'We then print the shot out on a sheet of paper, which traditional background artists use as a template to paint the detail of the object. All of the exterior detail and definition on the object is added at the painting stage – all the cracks, weathering, water stains and so on. We didn't want to use computer-generated lighting – which often looks too perfect – so all of the shadow and light information was added as 2-D detail in these paintings. When finished, they were scanned into the computer and then projected onto the model from the position that they were originally filmed from. Two or more projectors may have been used for each object to cover the detail of all of its sides.

'What we ended up with was the computer equivalent of the sound stage,' says Kraut. 'To build a street within the computer, we would position two rows of buildings, place another object underneath for the road, position a background painting in the distance and place a large dome with a sky painting mapped on it over the whole thing. The buildings with their projected surfaces looked like 2-D pictures, but we could animate our virtual camera to move past them, look around and do all the normal things that a live-action camera could. We wrote lots of software that would do things like blur the edges of the objects, so that they looked more than ever like 2-D paintings rather than hard-edged 3-D objects. For the first time ever, animation directors could actually create dynamic cinematography just like a live-action director, but their images looked like a camera was moving through a painting. If they didn't like a camera angle, the directors could just move the camera.'

Though the main characters in *The Prince of Egypt* were hand-drawn and animated in the traditional way, much of the film's additional cast of thousands was generated by computer. 'Descriptions of the Exodus in the Bible

actually mention 600,000 Hebrews. We didn't have quite as many as that, but we still had scenes with many thousands of people that could never have been achieved using traditional techniques,' says crowd animator Wendy Elwell.

'The first big crowd scenes are right at the beginning of the film, when we see hundreds of male slaves building a new temple complex,' explains Elwell. 'For these scenes we built a single 3-D digital character that matched the drawn characters. He was then reshaped to create a total of twenty different characters. These were then dressed with different hair, beards and clothing so that each person in the crowd looked more unique. We then animated walk cycles – sequences of movement that can be repeated as required – so that characters could walk for as long as was needed in any scene. Four separate cycles were used to give more variety of movement. The characters were then placed into scenes as 3-D objects, but they were rendered [236>] with flat colours and black outlines to make them look like hand-drawn 2-D characters. They were originally intended only to be in the background, but as the directors got used to using them, they put them nearer and nearer to the camera until their detail didn't hold up any longer. When the 3-D digital characters got too near the camera, they were automatically swapped with 2-D hand-drawn ones – though I defy anyone to spot where they change.'

Despite automated methods making their production more efficient than ever, the future for 2-D animated films currently seems somewhat uncertain. Disappointing box office returns for productions such as DreamWorks' *Spirit: Stallion of the Cimarron* (2002) and Disney's *Home on the Range* (2004) have been massively eclipsed by the popularity of 3-D animated films such as *Shrek* (2001) and *Madagascar* (2005). As a result, the major studios are now investing heavily in 3-D rather than 2-D production. Among the more successful 2-D productions of recent years have been the films of the Japanese animation director Hayao Miyazaki. His stylish and often stirring films, such as *Howl's Moving Castle* (2004) and the Oscar-winning *Spirited Away* (2001), have established a considerable following with Western audiences who might not normally consider watching a 'cartoon'.

3-D ANIMATION

The principle behind the animation of models is broadly the same as that which brings life to flat artwork, the only change being that it is three-dimensional objects rather than two-dimensional pictures that are moved fractionally and photographed frame by frame to produce the illusion of movement.

Dimensional animation was explored long before the advent of moving pictures. The popular zoetrope (<11) was not confined to giving displays of animated artwork. Rather than pictures, some zoetropes contained a series of carved models in varying poses – a spin of the wheel might send a wooden bird on a graceful flight around the tin drum, or make a small dog leap through a hoop. When photographic moving pictures arrived on the scene, the technology was quickly seized upon to breathe life into otherwise inanimate objects.

The fantasies of Georges Méliès (<14) were built around the ability to stop the camera, make alterations to the scene and commence filming again. Méliès used the 'arrêt' method (literally meaning the 'stop' method) to create many cinematic wonders. Producing the impression of motion by stopping the camera, moving an object or model and then starting the camera again has since become widely used and is known as 'stop-motion' animation.

As early as 1898, Méliès produced a short commercial in which wooden alphabet blocks danced around the screen and arranged themselves into the spelling of an advertiser's name. A year before, Albert E. Smith, a founding partner of the Vitagraph company, had enlisted the services of his daughter's wooden toys and, with partner J. Stuart Blackton, produced a short stop-motion film entitled *Humpty Dumpty Circus* (1897) in which the toys moved about the screen with a life of their own.


Perhaps due to the availability of ready-made objects for manipulation, the technique of stop-motion animation seems to have been considered especially suitable for the production of children's films. Biograph produced the animated *Dolls in Dreamland* (1907), featuring a troupe of dancing dolls and a sprightly teddy bear. The ever-inventive director Edwin S. Porter (<17) spent a week of 12-hour days creating the imaginative march of six teddy bears for his rather less imaginatively entitled film *The Teddy Bears* (1907).

There were also a number of early uses of stop-motion for films aimed at a more adult audience. The English producer Arthur Melbourne-Cooper animated a collection of matchsticks in a Boer War propaganda film called *Matches Appeal* (1899). Edwin S. Porter produced *The Dream of a Rarebit Fiend* (1906), in which nightmares induced by a late-night snack of cheese haunt a sleeper who envisages the contents of his bedroom dancing around him before the bed itself takes flight into the night skies.

As the live-action feature film and its narrative forms developed, stop-motion was used for short, usually humorous films and was rarely considered for anything more adventurous. However, in around 1914 a young man obsessed by the magic of bringing life to inanimate objects began to produce short animated films that would develop into the most spectacular and influential film fantasies ever seen. The early story of stop-motion in feature films is the story of Willis O'Brien.

LEFT: While 2-D animation faces an uncertain future in mainstream cinema, the beautifully rendered films of Japanese animator Hayao Miyazaki, such as *Howl's Moving Castle* (2004), have won a dedicated following and many awards.

BELOW: English producer Arthur Melbourne Cooper animated a collection of ready-made toys in a surprisingly realistic model town for his children's film *A Dream of Toyland* (1908).



PROFILE WILLIS O'BRIEN



Born in Oakland, California, the young Willis O'Brien (1886–1962) worked as a marble cutter, cowboy, prize fighter, and cartoonist for the San Francisco *Daily News* before being hired to create sculptures for the 1913 San Francisco World's Fair. While preparing some small clay figures for an exhibit on boxing, O'Brien carried out an experiment that would change both his life and the course of cinematic history.

Using a borrowed newsreel camera, O'Brien moved the clay pugilists a fraction at a time before exposing each frame of film. When the film was developed, O'Brien saw a jerky, spasmodic boxing match in miniature. The images were not perfect, but O'Brien had created life where there had been none, and he was gripped.

For his second experiment with what he came to call 'animation in depth', O'Brien produced a one-minute film featuring a caveman and a dinosaur. He spent the next two months filming in the basement of a San Francisco theatre to produce *The Dinosaur and the Missing Link* (1915), a five-minute comedy starring an apeman, cavemen and various prehistoric creatures, crudely fashioned from wooden skeletons with soft clay bodies. The film was so good that it was bought by the Edison film company for distribution, and O'Brien moved to the East coast with plans to produce more films under the banner of Mannikin Films Inc.

O'Brien's company produced a number of animated shorts for Edison between 1915 and 1918, with titles such as *Prehistoric Poultry* (1917) and *Curious Pets of Our Ancestors* (1917). Most important was a film called *Nippy's Nightmare*, in which O'Brien intercut his increasingly sophisticated dinosaur animation with separate shots of live actors in matching environments. This was perhaps the first time that a real person had co-starred with stop-motion animated creatures.

Hungry for better subjects and bigger films, O'Brien embarked on the production of a feature film, *The Ghost of Slumber Mountain* (1918), of which only 15 minutes survive. The animator worked with the American Museum of Natural History to produce dinosaurs that were thought to be scientifically accurate in appearance and behaviour. The film's poster declared: 'These giant monsters of the past are seen to breathe, to live again, to move and battle as they did at the dawn of life!' O'Brien even appeared in the film as an old hermit called Mad Dick. The picture was a great success and earned the attention of the film producer Watterson R. Rothacker, founder of the Industrial Motion Picture Company, which encouraged the development of special effects techniques.

Rothacker owned the rights to Sir Arthur Conan Doyle's novel *The Lost World* (published 1912) and on the strength of his previous efforts, O'Brien was hired to bring Conan Doyle's dinosaurs to life. Conscious that his dinosaur puppets – pliable clay bodies sculpted over wooden jointed skeletons – would not be sophisticated enough for *The Lost World*, O'Brien hired a young sculptor called Marcel Delgado, whom he had met at art classes.


Using the classic dinosaur paintings of Charles Knight for reference, Delgado spent two years coaxing 50 exquisitely detailed prehistoric beasts back from

extinction. Each 45 cm (18 in) creature was given an anatomically correct steel armature, complete with articulated spine and tiny ball-and-socket joints for every moving limb and digit. Once complete, the skeletons were painstakingly built up with cotton wadding and pieces of sponge to represent bulging muscle and flesh. Some models were fitted with air bladders that could be inflated and deflated during animation to mimic the animal's laboured breathing. The skeletons were supplemented with wire inserts where extra movement might be required, such as in the lips, eyebrows and cheekbones. Each body was then given an outer skin of liquid latex. Warts, scales and other surface textures were cast in latex and individually stuck onto the dinosaurs before being painted.

While Delgado laboured over the models, O'Brien oversaw the construction of the environments that they would inhabit. The *Lost World* itself was a miniature landscape, 60 x 90 m (200 x 300 ft), built at First National's Hollywood studios. The tabletop world had mountains, lakes and scaled-down tropical foliage with leaves made from sheet metal so that they would not grow, wilt or move in any way during the lengthy animation process.

By June 1922 O'Brien had completed a reel of animated dinosaur footage and it was shown to Conan Doyle, who was visiting the US on a lecture tour at the time. The author was delighted with the footage and could not resist a joke at the expense of his friend Harry Houdini. Houdini, who was highly sceptical about Doyle's belief in spiritualism, invited the author to the annual meeting of the Society of American Magicians in New York. At the meeting, Doyle screened O'Brien's dinosaur footage without explaining its origin. The audience was amazed and, incredibly, took what it saw to be genuine. The front page of the following morning's *New York Times* announced that prehistoric beasts had been discovered with the headline 'Dinosaurs Cavort in Film for Doyle'. Doyle immediately issued a statement revealing the true source of the footage. It was great publicity for the film and, above all, confirmed that O'Brien was producing extraordinary work.

With the effectiveness of his animation confirmed, O'Brien continued work on the film. He wanted to show something that had never been seen before: real people sharing the same space as animated creatures. This was achieved by building sections of full-scale set (usually trees or boulders) that matched the miniature set. The dinosaurs were then animated with a small area of the camera lens masked out (usually a lower corner), so that a portion of the frame was left unexposed. When the animation was complete, the film was rewound and a counter-matte placed over the lens so that only the unexposed area of the film was left uncovered. The film was then passed through the camera a second time to expose the performers in the full-scale set, apparently reacting to the prehistoric beasts in the rest of the frame. This static split-screen process was used throughout the film to combine not only actors and dinosaurs, but also dinosaurs and real locations. In the scene in which a



tyrannosaurus fights a triceratops, the matted-in Los Angeles river flows past in the foreground, helping to lend a sense of scale.

Not content with putting people into miniature scenes, O'Brien decided to place an animated brontosaurus into full-scale live-action scenes. Two thousand extras were filmed reacting to an invisible dinosaur in a London street set. The model creature was then animated in front of a plain white backdrop using a camera perspective matching that used for the live-action footage. Using the Williams process (<58) to create a travelling matte, the dinosaur was combined with the live-action plate. The result was convincing enough for contemporary audiences, though the heavy matte lines around the dinosaur seem obvious now. A full-scale tail and foot were also built for actors to interact with in these scenes.

The Lost World was a huge success when released in 1925; audiences flocked to see the two-hour, ten-reel film with its realistic scenes of dinosaur life. O'Brien's

animation was a triumph, not only as a technical achievement but for the subtlety with which his creatures performed. O'Brien gave his creatures personality and filled their performances with detailed character nuances, in the way they move, flail their tails or prepare to pounce. Tragically the film only survives in fragmentary form, much of the footage having been lost or destroyed since its release. An effort to restore the film has had some success, and a number of missing sequences have been rediscovered, restored and reintroduced for modern audiences to enjoy.

Despite the success of *The Lost World*, O'Brien did not work again immediately – the film industry had found a new gimmick in sound, and studios were more interested in hearing the voices of Broadway stars than seeing stampedes of dinosaurs. However, in 1930 O'Brien persuaded RKO to let him begin work on a film called 'Creation', which he believed would be the greatest dinosaur film of all (184>).

KING KONG

During the depression of the early 1930s, RKO hired Paramount producer David O. Selznick (1902–65) to help the company stave off bankruptcy. Selznick's assistant in this task was Merian C. Cooper (1893–1973), who, with Ernest Shoedsack (1893–1979), had made a series of successful adventure films including *The Four Feathers* (1929). One of the first decisions made by Selznick and Cooper was to cancel Willis O'Brien's *Creation* project, which had already cost the studio \$100,000 and showed no signs of nearing completion.

However, Cooper believed that O'Brien might be the only person who could bring life to a story that had been gestating in his mind for some time. After O'Brien produced some successful test footage, the project was green-lit under the working title of 'Production 601', later becoming known as 'The Beast' and then 'The Eighth Wonder'. Eventually, the film was revealed to the world as *King Kong* (1933).


For *The Lost World* (1925), O'Brien had built a large set that could be photographed from all angles. For Kong he created another exotic world for the film's star to inhabit. Before starting work on the film O'Brien had discovered the engravings of the French artist Gustave Doré, whose illustrations of forest scenes with shadowy fringes and scattered pools of hazy sunlight had a prehistoric feeling of impenetrable depth. O'Brien decided to re-create this look by building a number of small sets, each designed to look like a Doré engraving come to life.

The miniature jungle sets were filled with gnarled trees made from modelling clay and palm fronds fashioned from sheet metal. Each tabletop set was arranged in a number of planes, one behind the other. Each plane was separated by a sheet of glass on which Doré-style scenery had been painted by production artists Mario Larrinaga and Bryon Crabbe. Once properly aligned and lit, each forest scene gave the impression of great depth. O'Brien's multiplane model system was the stop-motion precursor of the system used to bring depth to the flat animation of Disney a few years later (<172).

Marcel Delgado was called upon to create the marvellous model creatures. The mighty Kong was portrayed by six 46 cm (18 in) puppets made using the same build-up technique that had given the creatures of *The Lost World* their convincing muscle and flesh. For Kong, Delgado even stretched rubber tendons between joints to give the ape a realistic



FIGURE 3 MINIATURE REAR PROJECTION



THIS PAGE: One of cinema's most iconic images: Kong battles for his life atop New York's Empire State Building.

BELOW LEFT: Kong battles with one of Marcel Delgado's beautifully modelled dinosaurs. The influence of Gustave Doré is clear in the multi-layered jungle backdrop.





LEFT: A performer cowers before a mighty monster. In fact, the performer and the tree are placed in front of a large rear-projection screen, and the dinosaur is an animated model just a few inches high.

BELOW RIGHT: Kong battles a pterosaur while clutching a puppet of the screaming Ann Darrow (Fay Wray). Such scenes used ingenious combinations of live action, stop-motion and rear projection to create their groundbreaking images.

sinewy appearance. Covering the ape's complex body was a patchwork of trimmed black rabbit fur. The steel armatures in Kong and his prehistoric co-stars were built with keyholes in the base of their feet. This allowed them to be secured firmly to the floor between takes using 'tie-down' pins screwed up through a grid of holes on the floor of the tabletop set.

Again, O'Brien needed to mix real-life actors with his animated wonders, but he wanted to avoid the process of double-exposed split-screen mattes that he had used in *The Lost World*, which was slow and carried with it the risk of ruining the original footage. The process of rear projection was being perfected in the early 30s and RKO was an early leader in the field, pioneering the use of a new type of background screen invented by paint department head Sydney Saunders. The first scenes to use the new Saunders screen were those of Fay Wray perched in a treetop while pre-animated footage of Kong fighting a tyrannosaurus is seen in the background. The new process took some time to perfect due to the difficulties in preventing foreground lighting spilling onto the background screen – one shot of Wray in the tree was achieved only after an exhausting 22-hour session.

Cooper asked O'Brien if the process could be reversed in order to place rear-projected actors into the miniature tabletop scenery. After much experimentation, O'Brien and his team perfected a method of miniature rear projection (fig. 3, <184). The model jungle settings in which actors needed to appear were built with small areas where images could be rear-projected onto miniature screens made from stretched surgical rubber. Before each shot was filmed, many tests were made in order to match the lighting of the miniature set with the exposure of the rear-projected images. During animation, the model creatures were manipulated fractionally, the rear-projected image was advanced by one frame, and a single frame of the composite image photographed before the process was repeated.

One of the first scenes to use the miniature rear-projection process was that in which Bruce Cabot (John Driscoll) hides from Kong in a cave. A miniature cliff was built with a small cave in its wall, and a rear-projection screen was set just inside the cave. Footage of the actor performing in a large-scale cave set was projected onto the screen one frame at a time, while the model of Kong was animated rampaging on the cliff above. This shot was later cut together with a close-up of the performer interacting with a giant arm that reached into the cave.

The animation of King Kong took 55 weeks to complete. Its key actors were called back to the studio every few weeks to film scenes in which they had to react to recently completed animation, or provide the performances needed for the next batch of miniature rear-projection shots. Some animation sessions turned into marathon events; the animation team quickly learned that shots begun one day and completed the next often took on awkward changes of pace or style mid-scene; once started, the animation of a shot therefore continued until it was finished.

There were other unexpected hitches in the animation process. On one occasion an animator was halfway through a shot when he noticed that a pair of pliers left on the set was just visible in the bottom of the frame. Not wishing to start the sequence afresh, the animator slowly animated the out-of-focus grey shape out of the shot, hoping it would look like a passing jungle creature. In another scene a primrose planted as part of the jungle foliage chose the day of filming to come into bloom. No one noticed the flower's cautious emergence during animation, but when the finished shot was viewed, the scene's prehistoric star was upstaged by the energetic emergence of a giant flower, wasting hours of work.

While the team of animators achieved much of the animation of Kong in long shot, O'Brien animated particularly emotional scenes and close-ups himself. As a result, Kong remains one of the most emotive creatures to have been created for the screen. Even Kong's often criticized bristling fur coat – caused by disturbance of the stiff rabbit fur during animation – seems to add to the great ape's personality.

King Kong was a sensational hit on its release in 1933 and again when re-released in 1952. The public flocked to see the film, thrilled both by the story and the wonders of a lifelike giant ape. Kong was undoubtedly the most extraordinary technical achievement of its time and a catalogue of the most modern special effects techniques, including animation, rear projection, miniature rear-projection travelling mattes and matte paintings, as well as clever optical work by Linwood Dunn (<72).

It is testimony to the remarkable and emotive personality of O'Brien's Kong that a number of remakes and imitations of this landmark film have been attempted. Despite the technical brilliance and extraordinary realism of Kong's most recent reincarnation (2005), however, his original performance retains a power and innocence that thrills to this day.



PROFILE RAY HARRYHAUSEN



In 1933 the 13-year-old Ray Harryhausen (1920–) watched the newly released *King Kong* at Grauman's Chinese Theater on Hollywood Boulevard. 'This ape was so huge, and there were so many awesome dinosaurs as well,' he remembers with glee. 'I've never been the same since I walked out of that theatre!'

The young Harryhausen became obsessed with *Kong* and with discovering the magic behind its creation. 'There had been other ape films before *King Kong*, where the gorillas were played by men in suits. I knew for sure that *Kong* was different, but I didn't know how. It took me over a year to figure out the glories of animation, and eventually I found a magazine that had an article that explained all about the stop-frame method.'

On discovering the technique, Harryhausen began to experiment with animation on his family's porch and later in a small studio built by his father in the garage. One day Harryhausen plucked up the courage to phone his hero, Willis O'Brien. 'O'Bie [O'Brien's nickname] kindly invited me over to MGM where he was preparing to shoot *War Eagles* – a film that never got made because of the war. I went to the studio and took along some of the dinosaurs I'd been making. He looked at my stegosaurus, which had won second prize in a competition at the local museum, and he said: "Its legs look like sausages. You need to learn about anatomy and how muscles work." So then I began to study anatomy, life drawing, sculpture and photography.'

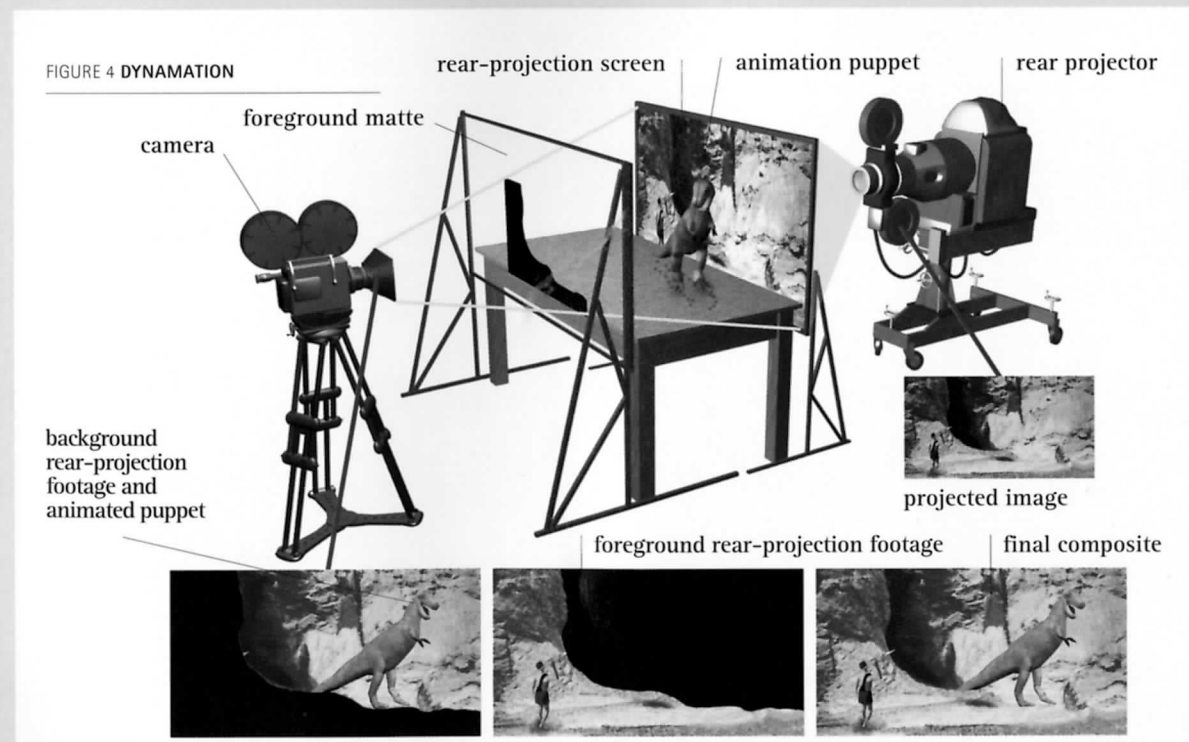
Harryhausen's animation skills quickly developed, and in 1940 he gained employment as an animator on George Pal's Puppetoon series (191>). After spending the war as an assistant cameraman and animator in Frank Capra's special-services film unit, Harryhausen began to

make his own series of animated fairy tales. Then, in 1947, having kept in touch with his mentor, Harryhausen accepted a job as O'Brien's assistant animator on *Mighty Joe Young* (1949).

'Working with O'Bie was a dream come true,' recalls Harryhausen. 'I helped during the whole pre-production phase, and when we got the go-ahead from the studio, RKO, I began work on the animation. I eventually did about 80 per cent of the character animation on *Joe Young*. O'Bie spent much of his time supervising all the other elements.'


The production of *Mighty Joe Young* used many of the techniques developed for *King Kong*. One small but significant advance was the use of rubberized fur for the gorilla puppets. Kong had been covered with rabbit fur, but the finer hair on unborn calf hide offered a better sense of scale for the smaller, 30 cm (12 in) Joe puppet. In a process invented by the taxidermist George Lofgren, the fur was embedded in paraffin, leaving only the skin exposed. The skin was then removed with acid and replaced with liquid latex. The finished fur, with its roots in latex, sprang back into place when touched during animation. Kong's bristling fur, which had annoyed some yet endeared him to many more, was not a feature of Joe's character.

After the release of *Mighty Joe Young*, Harryhausen was asked to be chief animator on the low-budget monster movie *The Beast from 20,000 Fathoms* (1953). 'Beast had a ridiculously tiny budget [\$150,000] considering its subject matter,' remembers Harryhausen. 'Doing things the O'Bie way was out of the question. He used big painted sheets of glass in several planes that gave a wonderful atmosphere but which were expensive



RIGHT: *The Beast from 20,000 Fathoms* (1953) was Harryhausen's first major film as chief animator. For scenes such as this, he devised a method of sandwiching stop-motion characters into live action using a split-screen rear-projection process that would become the basis for all of his later work.

BELOW: This sequence of images shows how Harryhausen combined rear-projected live action (a) with animation (b) to produce convincing composites (c). The final image (d) is a rare glimpse inside Harryhausen's studio during the filming of *The Valley of Gwangi* (1969). The model balloon and cage containing the dinosaur puppet are placed in front of a rear-projection screen. The wooden leg to the right of the image is part of the frame on which the animator placed mattes to conceal parts of the image during filming.



and restricted the camera to a single view. Also, the glass could crack in the heat of the lights. For *Beast* I devised a way of combining live-action actors with animated creatures and miniature backgrounds. It was the basis of the system that I used for the rest of my working life.'

The system that Harryhausen devised, which became known as 'Dynamation' for publicity purposes, was in many ways the opposite of that favoured by O'Brien. Rather than placing rear-projected actors into a miniature environment, Harryhausen used rear projection to place his animated characters into real environments containing real people. During live-action photography, performers reacted to the invisible creatures, sometimes looking at a cardboard head on a long stick as a guide. Harryhausen then studied the developed footage to ascertain what movements the animated creature should make in relation to the actors and their environment.


In his studio, Harryhausen then prepared an animation table (fig. 4) with a floor contoured to match the one that the creature was to traverse in the live action. A static split-screen matte that matched this contour was fitted to the bolted-down animation camera and used to mask out the bottom half of the frame, containing the animation table and other paraphernalia. The image exposed to the camera showed the puppet down to foot level, with the top portion of the rear-projected live-action background plate behind it.

Animation then proceeded as usual. Surface gauges (positional metal rods on stands) were placed at important points around the puppet's body to help mark the last position of its nose, tail, feet, and so on. The puppet's new pose was then created in relation to the action in the rear-projected frame of film and to the old positions marked by the gauges. When the new position was satisfactory, the

gauges were removed, the animation lights switched on and a frame exposed. The surface gauges were then brought back in, the rear-projected footage was advanced one frame and the next pose was created.

When the whole sequence had been animated, the puppet and animation table were removed from in front of the rear-projection screen. The camera negative and rear-projection footage were rewound to the start frame, and a counter-matte was applied to the camera to mask out the top section of the frame that had already been exposed. The camera negative was then exposed to the lower section of the background plate, so the two halves were married along the matte line. The final developed negative was a seamless composite of live-action footage and animated creature.

Harryhausen also made changes to the type of puppet he used for animation. 'Delgado's build-up method was terrific,' says Harryhausen, 'but it was time-consuming, and if a breakdown occurred during filming, the whole puppet might have to be stripped down and rebuilt.' Instead, Harryhausen sculpted his creatures in clay. These clay shapes were covered in plaster to create a mould. The creature's ball-and-socket armature was placed inside the mould, which was then filled with liquid foam rubber. When baked, the rubber set. The rubber creature was then painted and dressed with eyes, clothing, hair and props. If a mishap occurred during filming, the puppet could be stripped off and an identical rubber model quickly made. Harryhausen did not rely on this method to produce all of his creatures, however. For *Mysterious Island* (1961), a giant crustacean was made by fitting an armature inside a real crab shell. Another notable shortcut was to give the giant octopus from *It Came from Beneath the Sea* (1955)



ABOVE: *Jason and the Argonauts* (1963) is perhaps Ray Harryhausen's best-remembered film. This sequence, in which the Argonauts battle with animated skeletons, is certainly one of his most extraordinary achievements.

just six arms instead of eight. 'No one could have counted all those writhing arms, and it saved a lot of animating!' chuckles Harryhausen.

After a series of 'monster on the rampage' films in the 1950s, Harryhausen turned to the subject matter for which he became best-known. 'I had animated monsters destroying LA, Washington, San Francisco and New York – and there's a limit to the number of times you can do that sort of thing,' he explains. 'Charles Schnee [Harryhausen's long-term producer] and I decided that we wanted to move on.'

The 7th Voyage of Sinbad (1958) became the first in a series of films based on the *Arabian Nights* tales and Greek myths. Harryhausen considers this work to be his best, partly because stop-motion was so suited to the subject matter. 'I feel mythological stories are best served by stop-motion because the process produces a sort of dreamlike quality. It's true that the animation is not very realistic, but we never tried to mimic reality. We played on the melodramatic aspects of film, and Greek mythology is very melodramatic.'

Harryhausen's work in these mythological films was technically and artistically stunning, with the animator paying more attention than ever to the subtleties of characterization in his creations. 'I always tried to give my characters little habits that made them seem more believable,' he explains. 'It doesn't take much, just small habits such as taking a quick look at the ground before they step forward. I would also consider a creature's physique when planning their movement. In *The Golden Voyage of Sinbad* [1974], Kali moves in an unwieldy way because she is so top-heavy, and Talos in *Jason and the Argonauts* [1963] was actually criticized by one journalist for being jerkily animated – but he was designed to move that way because he's a giant metal statue with rusting joints!'

In the early 1960s, Harryhausen and Schnee moved their operations to England, which was closer to the

European locations that they often used and allowed them access to the sodium vapour travelling matte process (<64), which only Disney was licensed to use in the United States. During this period Harryhausen developed many new techniques for combining live action and animation more realistically. 'I tried all sorts of methods to make the two separately filmed elements look as if they were filmed together,' he says. One of the animator's cleverest deceptions involved using objects that appeared to cross the boundaries between live action and animation.

Typically, a live-action character might throw a spear through the air and it would appear to stick in the side of an animated creature. Such shots were achieved by animating the creature in front of rear-projected live action in the usual manner. As the rear-projected actor threw the spear, the weapon would, of course, disappear behind the model creature rather than stick into it. As the rear-projected live-action spear began to disappear behind the model, Harryhausen hung a miniature spear inside the model set so that, from the camera's viewpoint, the model spear covered the image of the rear-projected spear. As the live-action spear disappeared behind the creature, Harryhausen animated its miniature replacement so that it impaled the rubber model, producing the illusion that both live-action and animated elements were sharing the same space.

A different method was used to create the illusion that live-action characters with swords were interacting with animated creatures. Harryhausen mounted a sheet of glass in front of the puppet so that as a character thrust a sword at the creature, the tip (which in reality disappeared behind the model) could be painted onto the foreground glass frame by frame.

One of Harryhausen's most memorable and technically exacting achievements is the famous skeleton fight in *Jason and the Argonauts*. 'For that sequence I had to plan the movements in every single frame meticulously in order to animate seven skeletons simultaneously. Each skeleton had five appendages, so this meant I had to animate 35 separate movements for each frame, and each movement had to synchronize perfectly with the movements of the three live-action men.' It is hardly surprising that Harryhausen averaged just 13 frames a day during the four-and-a-half months that the short sequence took to complete.

Such dedication could be physically and mentally punishing, especially as Harryhausen was personally responsible for every frame of animation in all of his works, with the exception of his last film, *Clash of the Titans* (1981), when Jim Danforth and Steve Archer helped with some scenes. Extraordinarily, the master magician never received an Oscar for any of his films. In 1992, however, he received a special Academy Award in recognition of his lifetime's work. He also has a star on Hollywood Boulevard's Walk of Fame.

To this day, the methods used to create some of Harryhausen's best shots remain a secret. 'I never give everything away,' he muses. 'When a magician gives away all of his secrets, no one is interested any more!' Whether the technique behind his art remains a secret or not, the films of Ray Harryhausen continue to beguile.

REPLACEMENT ANIMATION

Most stop-motion animation is produced using a method called displacement animation, in which flexible models are moved fractionally between exposures. An alternative method involves substituting the entire model, or parts of it, between exposures. This rarely used technique is called replacement animation.

George Pal is often credited with pioneering the replacement method of stop-motion animation. Pal used replacement legs for his Puppetoon stars to produce their walking cycles. Whenever characters were required to walk, rather than reposition a flexible puppet between exposures, the entire bottom half of the puppet was replaced with one of a sequential set of carved wooden legs. Each set of around thirteen pairs of legs formed a walk cycle that was unique to the puppet in question. In 1943 Pal received a Scientific Academy Award for his development of the Puppetoon animation system.

One of the most groundbreaking modern uses of the replacement technique was for the stop-motion animated feature *Tim Burton's The Nightmare Before Christmas* (1993), directed by Henry Selick. The film used traditional ball-and-socket displacement armatures for the bodies of its main characters, which necessitated some 230 puppets in all. However, because the characters were required to talk, sing and express emotions beyond the range normally expected of stop-motion puppets, each of the major characters had a supply of replacement heads. A dialogue animator was employed to draw every conceivable combination of mouth and facial expression. The resulting 400 designs were then sculpted in modelling clay, moulded in rubber and cast in polyurethane plastic resin. Each head was then airbrushed to give it the correct colouring. One key character, Jack Skellington, required around 800 heads, allowing the expression of every possible emotion.

The finished heads were then photographed, digitized and stored in a computer databank. Each scene was then studied and the dialogue of every character broken down into frame-by-frame increments detailing the necessary phonemes and facial expressions. Using the computer database the various facial expressions were then assembled in the correct order to produce a video test of what the finished animation should look like. When approved, the computer performance was output on a breakdown sheet, telling the animators which heads to use in which order. During photography, animators first manipulated the ball-and-socket puppet body in the usual way before fixing on the necessary replacement head for each shot.



ABOVE: George Pal sits among some of the many replacement animation figures used for his Puppetoon series in the 40s.

BELOW: Computer-aided replacement animation was used to give life to the quirky characters in *Tim Burton's The Nightmare Before Christmas* (1993).



MOTION BLUR

Successful stop-motion animation can breathe the illusion of life into inanimate objects, but the method has a serious flaw that can prevent even the greatest animators from producing completely lifelike images.

When a real moving object is filmed at the standard rate of 24 frames per second, the shutter of the movie camera is actually open *during* the movement of the object. Since the object being filmed is in motion while its image is being exposed onto photographic film, the result is a photograph that contains a degree of 'motion blur' – a visible blurring that follows the most extreme movements of the object. Motion blur helps to impart the illusion of smooth, realistic movement when the still images are later projected onto a screen at rapid speed.

However, stop-motion animation creates the illusion of movement by projecting images of *still* objects in rapid succession. An animated object is moved *before* it is photographed, so its movements have no motion blur. The object's incremental yet distinct movements appear to jump from one position to the next during projection instead of flowing like real live-action photography. This jerky movement, which is called 'strobbing', is most pronounced in fast-moving animated objects that change position substantially from frame to frame. It is also particularly noticeable when animated objects that have no motion blur are combined with live-action actors that do have motion blur.

Many people believe that it is the very absence of motion blur that gives stop-motion animation its unique, almost magical quality. Without motion blur, the mythical monsters animated by Ray Harryhausen certainly have an ethereal presence that makes them particularly effective in the fantasy films in which they appear. However, for many productions absolute realism is the goal, and without motion blur, stop-motion animation can never be mistaken for live action.

The animator Jim Danforth used a time-consuming method to create artificial motion blur when animating a number of dinosaur sequences for *When Dinosaurs Ruled the Earth* (1970). The method involves exposing each frame of stop-motion animation three times. The puppet is moved into position and photographed at one third of the exposure normally needed to capture a good image. After moving the subject again, the animator exposes a second image onto the same frame of film, again at one third of the correct exposure. After repeating this operation a third time, the result is a single correctly exposed frame of film in which the animated object appears in three positions. The fast-moving parts of the subject – such as its legs – look partially transparent, because in these areas the three exposures record images of background scenery either before or after the object was moved in front of it. Though effective, this method requires the artist to animate three poses per frame of film, each pose covering one-third of the movement normally needed for one frame of animation. Rather than 24 separate movements per second of finished film, the process requires 72 moves, making it extremely time-consuming and impractical.

Another method of creating an artificial area of blur around an object is to smear petroleum jelly onto a sheet of glass placed between the camera and the object being animated. This process is also fairly time-consuming, because the grease regularly has to be cleaned from the glass and reapplied to match the motion of the object. Though not true motion blur, the unfocused edge around objects that this method produces can be quite convincing. The technique was used by Peter Kleinow to bring life to the Terminator robot in his animated sequences for *The Terminator* (1984).



ABOVE: The charm and humour of Nick Park's characters Wallace and Gromit have helped to ensure a place for stop-motion in modern cinema.

LEFT: This still from Ray Harryhausen's *Sinbad and the Eye of the Tiger* (1977) shows one of the flaws inherent in traditional stop-motion animation. While the live-action characters are blurred due to their movement, the leaping sabre-toothed cat remains completely sharp.


RIGHT: Phil Tippett admires the complex Go-Motion device built by ILM for the production of *Dragonslayer* (1981).

TOP RIGHT: While still requiring the consummate skill of highly talented animators, modern stop-motion films rely on digital technology to help create their imagery. After animation, this shot from Tim Burton's *Corpse Bride* (2005) was finished off by London's Moving Picture Company. Support rods were digitally removed from the main characters, background scenery was added, an otherwise difficult to animate veil was created digitally, and CG birds were made to flutter overhead.

Motion blur can also be produced in stop-motion animation by physically moving a puppet during exposure. This method was used by Phil Tippett (1942-) while animating the ED-209 robot for *RoboCop* (1987). After moving the robot puppet to each new position, a single frame was filmed in the usual manner. However, while the camera shutter was open and the film exposed, the model was physically 'wobbled' very slightly. Moving the model during exposure added a small amount of motion blur to the image, resulting in a very convincing animation.

Before Phil Tippett worked on *RoboCop*, he had already been part of an experiment that had culminated in some of the most effective model animation ever produced. During the making of *Star Wars* (1977), Industrial Light and Magic had pioneered the use of computer-operated motion-control cameras (<150) to film spaceship models one frame at a time while the camera moved past them. Since the camera moved while its shutter was open, the result was fast-moving spacecraft with realistic motion blur.

For *The Empire Strikes Back* (1980), the ILM team tried to apply this technology to the stop-motion animation of the woolly ice-creatures called 'tauntauns'. The tauntaun puppets were fixed to a motion-control mechanism that could move them backwards and forwards and up and down. The finer actions of the running tauntaun, such as its head and leg movements, were animated by Tippett in the usual fashion, but during photography, a motion-control mechanism moved the model vertically and horizontally while the shutter of the camera was open. The result was a form of motion blur that helped to eliminate some of the problems of traditional stop-motion.



Soon after working on *The Empire Strikes Back*, the ILM animation team were able to further develop their techniques for the production *Dragonslayer* (1981). The motion-control technique that had been used to animate *Empire's* puppet tauntauns was developed into a system that could control specific joints on the film's flying dragon. This system involved computer-controlled gears and motors that moved six rods. These rods were attached to key joints on the puppet dragon. The other ends of the rods were attached to a motion-control unit that could travel backwards and forwards on a 2.4 m (8 ft) track. The whole contraption was connected to a computer that could record and play back 19 channels of movement.

To produce the animated dragon sequences, the puppet dragon was animated in the normal fashion. Rather than photographing the dragon when the animator was happy with its pose, the exact position of the rods during each frame was recorded by the computer. Once the whole sequence

was programmed to the satisfaction of the animator, it was played back, the dragon mechanically repeating the movements given to it earlier. During playback, the dragon puppet's performance was filmed, its movements actually occurring *while* the shutter of the camera was open. The result was a dragon that moved with an astonishing fluidity and realism, thanks to the presence of real motion blur. This computer-age refinement of stop-motion was named 'Go-Motion'.

Though incredibly effective, Go-Motion is a time-consuming and expensive method and has only been used occasionally, perhaps most memorably to create the flying bicycles in Steven Spielberg's *E.T. the Extra-Terrestrial* (1982).

Today, computer-generated animation has largely superseded stop-motion animation as a method of creating naturalistic characters for the movies. However, modern audiences still enjoy the quirky immediacy of movies created entirely using stop-motion, though their painstakingly slow production means such films are few and far between. Aardman has had great success with its beautifully crafted crowd-pleasers *Chicken Run* (2000) and *The Curse of the Were-Rabbit* (2005), while Tim Burton built on the achievements of *The Nightmare Before Christmas* with his elegantly macabre *Corpse Bride* (2005). Though such films are still created using predominantly traditional techniques, their makers also now call on the latest technology to bring their visions to the screen.

Replacing the large film cameras normally used for stop-motion, *Corpse Bride* was shot using ordinary digital SLR stills cameras. London's Moving Picture Company (MPC) then used extensive digital paint techniques to erase the rods and wires used to support puppets as well as the complex motion-control rigs that moved the cameras during animation. They also added computer-generated elements that were carefully designed to look as if they had been traditionally hand-animated. These included the bride's flowing silk veil which would have been extremely difficult to animate with stop-motion. Environmental elements such as smoke, fog and water which are hard to create effectively using traditional methods were also produced digitally.

While modern stop-motion films still depend on the skills of talented animators, digital techniques now make production quicker and easier, helping to create the more sophisticated action and imagery that is expected by modern audiences.

PROFILE PHIL TIPPETT



After seeing *The 7th Voyage of Sinbad* (1958), Phil Tippett (1951–) became obsessed with animation. At the age of 13 he bought an 8 mm cine camera with money earned from mowing lawns and began to teach himself the art of stop-motion animation.

While still studying art at the University of California, Tippett gained experience animating a number of popular commercial characters such as the Pillsbury Doughboy and the Jolly Green Giant. Tippett's big break came when he joined the production team of *Star Wars* (1977), where he worked on alien designs and animated the holographic chess game between Chewbacca and R2-D2.

Tippett became a regular contributor to ILM projects and was part of the team that animated the AT-AT snow walkers for *The Empire Strikes Back* (1980). He also animated the film's renowned taunt sequences and developed a new method of producing stop-motion animation with motion blur. Tippett helped to refine the process, named Go-Motion, with

breathtaking effect for *Dragonslayer* (1981). By 1982 Tippett was head of ILM's creature shop, where he designed and built characters for *Return of the Jedi* (1983), winning an Oscar for his efforts.

In 1983 the animator established his own company, Tippett Studio, to provide sophisticated stop-motion animation for movies such as *RoboCop* (1987) and *Honey, I Shrank the Kids* (1989). Tippett's work with ILM on *Jurassic Park* (1993), which led the animator to enter the fledgling world of computer animation, earned him a second Oscar.

Despite running a large studio, Tippett has still managed to closely oversee work that combines digital technology with the performance skills of the traditional stop-motion animator. His studio has produced some stunning examples of CG character animation for films including *Starship Troopers* (1997), *Hollow Man* (2000), *Evolution* (2001), *Hellboy* (2004) and *Charlotte's Web* (2006). Tippett also directed the feature film *Starship Troopers 2* (2004).

3-D CHARACTER ANIMATION

3-D animation for feature films is now almost exclusively created with the computer. Like conventional stop-motion animation, 3-D computer animation requires the construction of models (albeit digital ones) that can be manipulated to produce moving images. 'Hard-surface' objects, such as spaceships and cities, are usually created using a selection of geometric objects that are combined and altered using a variety of modelling processes as described in the previous chapter (<156). CG models of 'organic' objects, such as human beings and creatures, however, require the creation of highly detailed forms that bear little resemblance to geometric building blocks like cubes and spheres.

There are several stages to creating a CG character model. Each stage can be achieved in a number of ways depending on the complexity of the character, how it will be required to perform, and the equipment and software used by the effects company producing the work. What follows is a broad overview of how CG characters are created and animated.

CHARACTER MODELLING

No matter how a CG character is to be produced, it will always start life as a series of drawings and paintings produced by a film's art department or concept artist. Paper designs approved by the movie's director will usually be turned into a series of small clay sculptures called 'maquettes' which allow a character to be studied from all angles before any further development.

When the final design of a character has been agreed, a much larger, fully detailed version will be made. Final character models are usually sculpted in a material such as Sculpey, a polymer modelling clay that is baked solid to create a permanent reference model. This final model is normally made in a neutral 'T' pose with the limbs fully extended to allow access to every area of the body. Complicated characters are sometimes sculpted in several separate pieces – such as body, head, and limbs – and these are only stitched together to form a complete character once all the parts are finally in the computer. A character's face will usually contain much more subtle detail than the rest of its body and so facial sculptures are often produced at a larger scale. To save time only half of a face or torso can be sculpted. When that half is in the computer a mirror-image copy can be made before being stitched back to the original to produce a whole, symmetrical character.

Once final character models are sculpted they need to be digitized – converted into digital models that can be further refined and animated in the computer. There are two common methods of turning a physical sculpture into a digital model: touch probes and laser scanners. It is also possible to extract 3-D information from photographs of sculptures using photogrammetric techniques (<160).

Touch probes were the first widely used method of converting sculptures into digital models. When using this method the topology of the digital model is first planned by hand-drawing a grid on the surface of the sculpture. The grid divides the surface of the model into various sizes of square.

BELOW: Here a touch probe is being used to digitize a sheep sculpture to create a digital model for *Babe 2: Pig in the City* (1998).

BELOW RIGHT: During a cyberscan (left), a laser beam scans the object or person to be digitized. The resulting 'cloud' of points can be greatly reduced without losing key detail (centre). The final model (right), once surfaced, has the scanned colour information mapped back onto it to produce a digital head.



Sparsely detailed areas that won't move very much, such as the back of someone's head, are covered by large squares, while delicate areas that will be extensively animated, like a face, are covered by many smaller squares.

The sculpture to be digitized is placed on a surface that emits a magnetic field. Connected to this surface is a digitizing probe – similar to a pen. The position of the probe's nib can be measured wherever it is placed within the magnetic field. Holding the probe, the modeller systematically works their way over the surface of the model, touching and recording every intersection on its surface grid. The result is a digital 'cloud' of points. These points are vertices [<154], and when they are connected together they form a mesh of polygons – a digital copy of the clay sculpture.

After digitizing, a mesh model needs to be 'cleaned'. This involves addressing the mesh to make sure there are enough polygons to produce the detail in areas where it is needed, but not so many that the model is too complicated to be used efficiently. Attention is also paid to complex areas, such as where ears join with a head or a tail meets a body, to make sure there are no complicated folds or intersecting polygons. At this stage a polygonal model might be converted into a NURBS model (<158) or it may have areas that require lots of subtle detail to be turned into subdivisional patches for further refinement (<157).

Touch-probe scanning has now largely been replaced by laser scanning. This method, also known as 'cyberscanning', uses a laser to optically measure the surface of an object and convert it into digital information. 'Laser scans are very good for capturing any object with very fine features, especially human faces,' explains Sean Varney, head of metrology at London's Framstore CFC (metrology being the science of measurement). 'The touch-probe process is completely impractical for capturing human features,' continues Varney. 'Most actors wouldn't be very happy about having a grid drawn on their face before being poked with a digitizing probe for an hour. They would also have to sit impossibly still within the magnetic field during the whole process. A laser scan, on the other hand, takes about twenty seconds.'



Scanners used for small objects, such as models and sculptures, have a fixed scanning head that shines its laser onto a small turntable. The object being scanned then rotates so that every part can be read by the laser. There are also hand-held scanners that resemble the devices used to read the bar codes on larger items at supermarket checkouts. These devices are swept over one small section of surface at a time. The information gathered in separate passes is then recognized by the computer and automatically 'stitched' together to form a complete model. Scanners used to digitize human faces and bodies are usually much larger and are mounted on a rotary arm that

swings the laser head itself in a complete circle around the subject.

'We can scan sculptures, props, animatronic characters, an entire human body, or even a live horse if you want,' says Varney. 'But most of the time it is just the face of an actor that is needed. In this case the subject to be scanned sits on the scanner platform with their head directly at the centre of the scanner's circular path. Before a scan, we give any hair on the head or face a dusting of cornflour. Hair can cause problems because it can work like thousands of prisms that refract the laser light in all directions. By giving hair a sprinkling of flour, more of the laser light is reflected back into the scanner to produce an accurate model. During scanning the scanner swings around the subject, shining a vertical beam of harmless laser light onto the face. The laser light is reflected back into two CCDs [charge-coupled devices (<93)] that convert the light into digital information. The scan records the physical contours of the face but also its colour information.'

As with the touch probe method, the result of a laser scan is a cloud of vertices that are linked to create a polygonal mesh. 'The scan collects a vast amount of information about the subject,' says Varney. 'In fact, the technique produces far more data than is needed for most purposes, so we edit the data according to how the model is going to be used. Automatic decimation software studies all the vertices and works out how to produce a highly detailed model with fewer vertices, then we finish it off by manually removing vertices or adding them back in where they are needed. A human head for use in a feature film usually starts life as a mesh of around one million polygons that can be edited down to about 10,000 polygons and still look very good. This produces a model with fewer polygons where they aren't needed, and more around areas that need to be flexible when animated – such as the corners of the mouth and eyes. Once the polygonal model has been finalized, the colour information that is also collected during the scan can be mapped back onto the model to create a convincing replica of a human head.'

According to Varney, laser scanning is becoming an increasingly popular safety tool for large-budget special effects productions. 'These days movies are altered so much during post-production that it pays to have a scan of all of the lead actors. If an important new shot is needed after filming has finished, it's possible to graft the scan of an actor's head onto the body of a double. Increasingly therefore, we are asked to scan the whole cast of a film,' says Varney. 'When we scanned the cast for *The Phantom Menace* [1999], we took our equipment to the studio and the performers had their scans done whenever they had a spare half-hour in their schedule. For bigger productions we also routinely scan props and even whole film sets so that alterations can be made long after filming.'



DIGITAL SCULPTING

Complicated organic forms have traditionally been difficult to create digitally, making it necessary to scan detailed clay sculptures in order to produce digital character models. However, a revolutionary software program called Z-Brush now means that intricate models can be sculpted directly within the computer without the need to make and scan sculptures in clay.

'Z-Brush has totally changed the way we make and animate digital characters,' claims computer graphics supervisor Vince Cirelli of Luma Pictures, who used Z-Brush to create characters and environments for *Underworld: Evolution* (2006). 'Z-Brush is exactly like having a lump of digital clay that can be sculpted on the computer monitor. A few years ago only computer experts could make sophisticated digital models, but now we can hire traditional artisans to sculpt directly in the computer.'

To produce a character using Z-Brush the artist begins by assembling a series of simple spheres and tubes to produce the rough, overall form of a character. This is exactly like throwing together lumps of real clay to make a basic sculpture – with a lump for the head, one for each leg and each foot, and so on. With the rough shape assembled, the computer then moulds it into a well-ordered polygonal mesh, quickly turning a crude collection of primitive shapes into something that already resembles a completed character.

The artist then starts to add more refined detail to the polygonal mesh, perhaps bulges for muscles or depressions for the ribcage. This is achieved using a series of paintbrushes that work like traditional sculpting tools, adding or removing clay where needed. To create a ribcage, the artist can simply paint onto the model with the appropriate size and strength of brush and an indentation is sculpted out of the surface. This happens in real time, as if clay was actually being carved out of a physical model. According to Vince Cirelli, very delicate detail can even be sculpted automatically, saving huge amounts of production time: 'In the past, if a sculptor wanted to create a lizard they would painstakingly scratch each

scale into the clay with a toothpick,' he says. 'Now they can take a photograph of a real lizard and use a clone brush to copy the scales in the photograph, producing actual sculpted scales on the surface of the model!'

As the model progresses the artist creates increasingly fine detail, moving from bold anatomical shapes through to delicate features such as wrinkles and skin pores. Each time the sculptor wants to add more subtle detail they can create a new resolution level, each new level quadrupling the number of polygons used to describe the surface of the model. If at any time more significant alterations are needed, the artist can go back to any of the previous levels of detail – even as far back as the original collection of spheres – and make that change. That alteration is then automatically updated in all existing levels of detail.

When sculpting is complete the result is an extremely refined polygonal model with an incredible level of detail. Such complex models would be slow and impractical to animate, so the final model is used to produce a black-and-white displacement map (<166), which is a 2-D record of all the delicate surface information.

'When we've made displacement maps we can take our final polygonal mesh and strip away almost all of the detail. A mesh with over a million polygons can be reduced to a mere 50,000,' explains Cirelli. 'We then use that low-polygon model for animation purposes. When we've finished animation the displacement maps are used to add back all the complex surface information during rendering, producing an incredibly detailed final image.'

Even the amount of detail created by a displacement map can be varied. 'For *Underworld: Evolution* we created an extremely complex castle model using Z-Brush,' says Cirelli. 'The displacement maps created the physical geometry of every single brick and tile during rendering. However, much of the time that extreme level of detail was not needed and would take an unnecessary amount of time to process. Therefore, we linked the level of surface displacement to the proximity of the camera to the building. If the camera was a long way off the walls remained relatively flat, but as the camera moved nearer the bricks literally "grew" outwards from the walls and became far more detailed.'



LEFT: Z-Brush was used by Luma Pictures to sculpt digital character models for *Underworld: Evolution* (2006). The character is sculpted with fine surface detail (top left). This detail is saved as a displacement map. Very subtle colour information can be painted onto the digital model (top right). Even a highly complex model can be reduced to a low-polygon mesh (bottom left). When surfaced, the basic mesh shows little detail (bottom centre). When displacement maps from the original sculpt are applied, the surface of the character becomes minutely detailed (bottom right).

TOP RIGHT: The final CG version of the vampire Marcus, as seen in *Underworld: Evolution*.



PROFILE JOHN LASSETER



Like many contemporary American animators, John Lasseter (1957–) was a graduate of the Disney Character Animation programme, run by the California Institute of the Arts in the 1970s.

While working on Disney features like *The Fox and the Hound* (1981), Lasseter saw the early computer animation being created for *Tron* (1982), and became fascinated by the potential of the computer. Lasseter then produced his own experimental film based on Maurice Sendak's book *Where the Wild Things Are*, showing how traditional hand-drawn animation could be combined with computerized camera movements and environments.

In 1984 Lasseter left Disney to join the computer graphics division of Lucasfilm, where he worked on a number of 3-D computer animation projects including the short film *The Adventures of André and Wally B.* (1984). He also animated the CG stained-glass knight for *Young Sherlock Holmes* (1985).

In 1986 Lucasfilm sold its graphics division to Steve Jobs and Pixar Animation Studios was born (198>).

At Pixar, Lasseter created a number of groundbreaking shorts including *Luxo Jr.* (1986) and *Tin Toy* (1988), which became the first CG film to win an Oscar.

Lasseter became widely praised for his extraordinary skill in creating strong characters from expressionless objects. In *Luxo Jr.*, for example, he created a touching father-and-son relationship between a pair of animated desk lamps in the space of just 90 seconds.

In 1991 Pixar signed a three-picture deal with Disney. The result was *Toy Story* (1995): co-written and directed by Lasseter, it was the world's first computer-generated feature film. It was a huge commercial and critical success and Lasseter received a Special Achievement Oscar. Lasseter has continued to oversee the creation of Pixar's innovative films and has directed *A Bug's Life* (1998), *Toy Story 2* (1999) and *Cars* (2006).

With Disney's acquisition of Pixar in 2006, John Lasseter was named chief creative officer, charged with overseeing all of Disney and Pixar's animated output and even contributing to the design of new attractions at Disney theme parks.

TOY STORY

Released in November 1995, *Toy Story* was the world's first completely computer-animated feature film, with over 70 minutes of 3-D digital animation.

Toy Story was created by Pixar, a company that began life as the digital research division of Lucasfilm. Under the leadership of Dr Edwin Catmull since 1979, Pixar created software used by ILM to provide effects for films such as *The Abyss* (1989) and *Jurassic Park* (1993). However, finding themselves more interested in producing original animation rather than technology for others to use, the Pixar division decided to split from Lucasfilm and in 1986 was bought out by Steve Jobs, co-founder of Apple Computers, for \$10 million.

After producing a number of groundbreaking short films, Pixar's aim became the creation of a full-length feature film. Working towards this goal, Pixar developed three core software systems: 'Marionette', a program for modelling, animating and lighting; 'Ringmaster', a program for scheduling and coordinating animation during production; and 'RenderMan', a sophisticated rendering program for producing vibrant, photorealistic images (237>).

In 1989 Pixar signed a three-picture deal with Disney and they began to plan their first feature. For their debut, Pixar decided to play to their strengths. Although their software was not yet capable of portraying convincing human characters, in his short films director John Lasseter (<197) had managed to elicit extraordinarily emotive performances from objects such as desk lamps and, in particular, toys.

The final script was storyboarded and assembled into an 'animatic' – a rough shot-by-shot version of the film with temporary dialogue laid over it. As production progressed, these elements were gradually replaced by low-resolution test animations, actual voice-overs and final rendered shots, slowly building up the finished film.

After the various characters, sets and props had been designed on paper, they were either built directly in the computer or were digitized from traditional clay sculptures. The computer models were then given varying numbers of control points that could be animated to produce performances. A complex character such as Woody had as many as 100 facial controls. To make life easier for the animators, some of these controls were clustered into basic movements and expressions – to create a smile the animator simply turned up the 'smile' control, for example. Performers such as Tom Hanks were videotaped during the recording of voice-overs and their movements and expressions were used by animators as a guide when creating the performance of their digital alter egos.

One of the complaints about early computer animation was that it looked 'plastic'. *Toy Story* proved that stylized yet natural-looking environments and characters could be produced and sustained. Every element of a scene was carefully painted to produce a detailed 'caricature' of real life. Even skirting boards were scuffed and carpets stained. It took 1,000 megabytes to store the movie's 400 models and 3,500 textures and the final rendering process took 500,000 machine hours – that is, 57 machine years.

Toy Story was a massive worldwide hit. Although much publicity was gained from its status as the world's first computer-generated film, reviewers and audiences alike appreciated the film for its exciting story and delightful characters. Pixar has since continued to create increasingly sophisticated and successful productions including *Finding Nemo* (2003), *The Incredibles* (2004) and *Cars* (2006). In January 2006 Disney acquired Pixar in a deal worth \$7.4 billion.



CHARACTER RIGGING

When the digital model of a character has been constructed, the result is a mesh of polygons or NURBS patches (<158) that form the character's outer shape. To make this empty skin suitable for animation, it must be filled with the digital equivalent of flesh and bones and then instructed how to move. This is a process called 'rigging'.

'At this studio the rigging department is called the Puppet Department,' explains Tippett Studio puppeteer Eric Jeffery, who has rigged characters for movies including *The Adventures of Sharkboy and Lavagirl in 3-D* (2005) and *Shaggy Dog* (2006). The term is a hangover from the stop-motion days when our puppets were painstakingly engineered with ball-and-socket joints and rubber muscles. Today we're essentially doing the same thing, but with digital models instead of physical ones.

'Before we start rigging a character we read the script, study storyboards and talk to the animators about how they plan to use the model,' says Jeffery. 'We need to know exactly what a character will be required to do. We also look at a lot of reference material and photographs of real animals. This helps us get a sense of a creature's anatomy and how the bones and muscles should work together to make the character move.'


When a new mesh is received from the modelling department, riggers first create a very simple skeleton that will allow the character to be tested by the animators. 'It takes us about a week to build the initial skeleton and get the character moving,' explains Jeffery. 'Then we give that low-resolution version to the animators, who produce a few sequences of the character walking, running and moving in extreme poses. This enables us to see if the mesh is functioning satisfactorily, whether it's capable of moving in the right way, or whether it needs any changes to the design. When we've made any alterations and are sure everyone is happy with the model we then pass that low-res rig back to the animators and they use it to begin producing animation. While they do that we start work to produce a fully rigged version.'

The full rigging of a character begins with the building of a refined skeletal structure. With the character positioned in a neutral pose (the classic 'T' stance in the case of bipeds), joints are placed inside the mesh and arranged exactly like those in a real skeleton. These joints can be programmed to rotate, twist, scale (change in length, height, or width) and 'translate' (move in any direction to a new position). 'Unless we're working on particularly flexible characters such as those we made for *Son of the Mask* (2005), we constrain each joint so that it can only perform in the way that it would in nature,' says Jeffery. 'That way the animators can be sure it will behave accurately when they create a performance. However, we always give the animators the ability to switch off these constraints – nature often has to be tweaked just a little to make a shot work.'

With the joints in place, riggers consider what type of skeletal structure to build. 'We have two approaches to building skeletons, depending on the type of character we're working on,' says Jeffery. 'If it's a principal character, or one on which we will be seeing a lot of bare flesh, such as our Abe Sapien character for *Hellboy* [2004], we tend to construct an anatomically accurate skeleton made up of

RIGHT: To create the incredibly lifelike aquatic character Abe Sapien for *Hellboy* (2004), Tippett Studio built up layers of anatomically accurate muscle which would be seen stretching and bulging beneath the character's skin. *Top*: The original background plate showing an empty tank of water into which the character would be placed. *Centre*: CG set-up showing muscled version of Abe. *Bottom*: Final composite, complete with watery atmosphere in front of CG character.

BELOW: This digital model of a horse was built by Tippett Studio. The first image shows the basic underlying skeleton and joints. Subsequent images show the anatomically accurate layers of bone, ligament and muscle that were layered over the top. The final low-resolution horse is what the animator sees when manipulating the character on their computer screen.



realistically shaped bone geometry. This will give us somewhere to fix muscles and will help to produce realistic movement on the surface of the character. However, if the character will only be seen in a few shots, or is perhaps covered in clothes or fur, such as our Templeton the rat for *Charlotte's Web* [2006], we won't bother with any bones and will connect the mesh directly to the joints.'

With joints and any bones in place, riggers next define the way they can move in relation to one another. A skeleton's joints are programmed to act in a hierarchical fashion, each joint forming part of a chain that is linked in a one-way relationship. A hierarchy begins with a parent object. Branching from that parent can be any number of child objects. Each of these 'children' can also act as 'parents' to further child objects, and so on. For example, the pelvis would be parent to the leg bones which would be parents to the feet which, in turn, are parents to the toes. When a parent object is moved during animation, it influences each of its child objects, but moving a child object does not necessarily affect the movement of its parent. Hierarchies are useful in many modelling and animation situations but are particularly appropriate when creating characters, since the skeletons of most creatures work in a hierarchical fashion. When a human being moves an upper arm, for example, the forearm, wrist, hand and fingers will all move in response.

'Skeletons end up consisting of a number of parent-child groupings,' elaborates Jeffery. 'These hierarchies branch backwards from the outer extremities to an overall parent object, which is usually located at the creature's centre of balance – typically somewhere a little above the pelvis in a bipedal character. The animators can move this central point around depending on the performance that is needed. For example, if a character is going to hang from its hands during one scene, their centre of balance will be much higher on the body; probably somewhere near the neck. The whole hierarchy will therefore be rearranged to link back to a parental point at the top of the spine, around which most of the character's motion will then swing.'

A hierarchical system in which the movement of parent objects affects the movement of child objects is called forward kinematics (FK). However, a more sophisticated hierarchical system called 'inverse kinematics' (IK) can be used to make many of the most complex movements of a model work 'automatically' during animation. Inverse kinematics uses a hierarchy that is engineered to work backwards – that is, the movement of a child object can affect the movement of its parents as an automatic function during animation. 'Inverse kinematics is a really great way of animating complicated movements such as walking,' says Jeffery. 'If you use forward kinematics to animate a walk cycle you have to individually move the upper leg, then the lower leg and then the foot, making sure at the end of the pose that the leg is bent in the correct walking position and that the foot is touching the ground appropriately. But if you then decide to change the leg position even slightly the foot will move in response and will need to be repositioned again. But with IK the animator can grab the feet and move them around and the parent legs will automatically follow, bending in the right way.'


Inverse kinematics is a useful tool that can save animators a great deal of time and frustration. However, some animators try to avoid such automation because they believe it can make character actions appear artificial and predetermined. 'We usually place both IK and FK skeletons inside our characters,' comments Jeffery. 'That way the animator can switch between the two to get the best motion for every movement a character makes.'

Once the skeleton of a character has been built, it is used as the framework for more internal construction. When Marcel Delgado built stop-motion puppets for Willis O'Brien (<182), he padded their metal skeletons with wads of cotton to represent the flesh and muscle beneath the skin. Digital character models are usually constructed in a similar way.

'Once the anatomical bones are set up we start to add the padding that goes on top,' says Jeffery. 'All the major muscles are replicated as meshes that have the appropriate elasticity and which will maintain their volume when they stretch or contract. The ends of the muscles are then attached to the bones in the correct places so that they flex, bulge and slide realistically when the joints move. We have to build the muscles up slowly, making sure that each new muscle works correctly and doesn't adversely affect those already in place.'



105
hbAS5 40 dlink



105
hbANIM1skin.29 pthuriot



105
hbAS5 89 jrm



Sometimes a muscle can get confused and pop inside out, flip back-to-front, or just break – messing up everything that is linked to it. Just like a real body it can be a pretty carefully balanced system.'

The complex construction of anatomically correct muscles is not always necessary for the creation of convincing characters. Those creatures whose skin surface will not be seen in detail are not normally given bones or muscles; instead, riggers use a technique called 'skin cluster weighting'. With this system the outer mesh of the character responds directly to the movement of the joints below. Several layers of greyscale maps are attached to the surface of the mesh, indicating the degree to which each area of skin is influenced by the underlying joints. When animated, the surface of the mesh will therefore move primarily in response to the motion of the nearest joint. It will also react, though to a lesser degree, to the combined movement of the other joints in the wider area. This system can sometimes produce unrealistic body forms, however. For example, a bending elbow joint can cause the surrounding mesh to lose its volume, resulting in an arm bend that looks more like a kink in a garden hose. In order to ensure that anatomically correct body shapes are maintained during a character's movement, the computer can detect what forms a mesh is tending towards and automatically place pre-sculpted 'blend shapes' (213>) in that area. The mesh will then blend in and out of these shapes over a period of several frames to produce the correct overall body form.

For models that do have a full system of anatomically modelled bones and muscles, the final additions to the body interior are areas of loose 'flesh' or 'fat' that are programmed to automatically jiggle with varying degrees of flexibility in response to the movement around them.

Finally, the outer mesh or skin of the model must be instructed how to move over the underlying biomechanics. 'The skin of all living creatures varies greatly on different parts of the body,' says Jeffery. 'With humans, for example, it can range from the very taut skin which is fitted quite closely over a kneecap to the taut yet loosely connected skin of the upper back which allows the shoulder blades to move freely beneath. Then there are flabby or loose areas of skin like those on the average person's

stomach. All these qualities need to be replicated in our more complex digital characters.'

The performance of the skin is driven by the motion of the underlying fat and muscle which is itself driven by the motion of the joints. 'We have to tell the skin exactly how to move in reaction to those various inner motions,' says Jeffery. 'Again, we paint a number of greyscale maps which instruct each area of skin how to operate. These maps can affect how thick the skin is and therefore the degree to which it might wrinkle during movement. They can also determine how stretchy the skin is when movement causes it to be pulled over a bulging muscle. They can even identify the points at which the skin is actually attached to the body and the degree to which it can move away from those points. This can be imagined as lots of bits of elastic tying the skin to the body at various points. The bits of elastic can be attached to different parts of the body and some pieces can be longer or tighter than others, thus affecting how the skin moves around during a character's motion.'

The final visual look of the skin, its colour and texture, is the result of hand-painted texture maps created in the studio's paint department (203>).

Once a character has been fitted out with bones, muscles and skin, riggers make the final preparations before animation. Their aim is to make the control of a character as simple and intuitive as possible for the animators. 'Each character will eventually have dozens of controls that the animators can use to produce a performance,' says Jeffery. 'The most fundamental are the handles used to push, pull, and place the limbs during animation. We build large boxes around the major body parts such as hands and feet so that they can be quickly grabbed and dragged around the screen just as if the animators were using a traditional puppet.'

Riggers also set up a range of automated actions that will help create a convincing character. 'The idea is that many subtle movements will occur as an automatic response to the broader performance created by the animators,' says Jeffery. 'For example, we will program the muscles in the legs and thighs to jiggle when the animator makes a character stamp its foot down to the ground. We will also create a breathing cycle in which the character's chest rises and expands, the shoulders move back and the stomach pulls in. This will be layered into the performance without the animator having to think about it. However, we do give animators the option of altering various parameters in order to change the speed of the breathing or make it heavier, should the action require it'. Control of all of a character's functions can be accessed through a series of on-screen pick-lists which sit alongside the puppet. These menus start off with the broader, most frequently needed groups of controls and are then subdivided down into increasingly subtle aspects that the artists may only occasionally want to manipulate during animation.

When the fully rigged character is complete it is ready to have the motion being produced by the animators applied to it. 'The full-size rigs are very large and complicated models that are impractical for use during animation,' says Jeffery. 'The animators use the much lower-resolution version that we made for them earlier in the process so that they can move it around on their monitor in real time and animate really quickly. This version often shows the character as little more than a few connected pieces of primitive geometry – boxes for bodies and spheres for heads – but it's enough for them to see how their motion is working. When they have finished animating or want to see how a performance is looking, their low-res puppet is temporarily attached to the high-res version. The low-res puppet then drives the motion of the high-res puppet and we can see how things are working.'

Once a character is handed over to the animators the job of the rigger is still not finished. 'We will continue working on a character right throughout production,' says Jeffery. 'The animators will always ask us to add new functions or features. Then the character might malfunction and we'll have to trace things back to the point where they went wrong – there's always quite a bit of maintenance to be done. Sometimes the director may even decide they want to change the look of the character altogether and we will have to rig a new mesh and merge the animation that has already been created with that new model. By the end of a show we may have been working on a character for over a year. We get to know them pretty well.'

CHARACTER PAINTING

One of the final tasks in the creation of a computer character is the painting of texture maps (<164) to add both colour and detail to the modelled exterior. Several methods can be used to create character texture maps, depending on the needs of the character and the techniques preferred by the visual effects studio.


A character's skin might be 'unwrapped' to produce UV maps – 2-D images that look as if the skin has been peeled off in sections and laid flat (<164). These flat maps are then digitally painted before being wrapped back onto the character. The main flaw of this system is that some warping or stretching of the 2-D texture is inevitable when it is wrapped back onto the 3-D model. Furthermore, painting a representation of 3-D detail onto a 2-D image is hardly an intuitive way for artists to work.

More commonly, characters are painted using a technique called 'projection painting'. This process allows artists to see painted detail appearing directly on the surface of the model as they work. To paint a character in this way, the artist first views the model from one fixed angle on the monitor – the side profile of a face, for example. They then start to paint over this view using digital brushes and paint. The artist is not actually painting onto the 3-D surface of the model, rather the digital paint is being applied to a sheet of virtual 'glass' that is placed just in front of the model. This sheet of glass works somewhat like a slide in a projector – throwing the painted image onto the surface of the model. When the artist is satisfied with how the colour and texture look from this angle, the model will be turned around so that another view is presented – the front of the face, for example. As the model is rotated, the first painted texture map moves with it, still being projected onto the side profile for which it was designed.

The front view of the face will not yet have had any paint applied to it. However, some areas will display detail that is spilling over from that being projected onto the side profile. This is just like some of the image from a projected slide missing the edge of the screen and falling, distorted and out of focus, onto objects in the background. In the case of a model head, some of the colour painted along the edge of the nose in the side profile will probably now be streaking along the cheek. The artist paints the front view of the face on a new sheet of glass, covering up any of the distorted detail that has spilled over from the side-profile projection. This process continues around the model, using a new piece of glass for every angle painted, each time covering any spillage from the previous projection. A human head would typically be painted in eight equal sections. The result is eight 'stacked' layers of glass which can be collapsed down into a single layer – the top layers permanently covering any of the spilled and distorted paint from the lower layers. This single final image looks distorted and stretched, but once projected back onto the model it creates a perfect surface that can be viewed from any angle.

Projection mapping can also be used when creating a digital version of an actor for use in shots where a computer-generated character is used to perform stunts, or a real stunt performer has their head replaced by that of the star. In such cases the star will be photographed from numerous angles and projection painting used to make each projected photograph blend seamlessly with the next on the surface of a cyberscanned (<195) model head.

As with other texturing processes (<164), colour is just one of the attributes that need to be created for the exterior surface of a character. Artists will also paint black-and-white bump and displacement maps (<165) as well as maps able to affect the skin's opacity, specularity, reflectivity or iridescence. These will all be used in conjunction with shaders (<167) to produce the final look of the skin.




TOP: The finished, untextured model of Abe Sapien in *Hellboy* (2004), showing the sculptural detail.

ABOVE: The final version of Abe Sapien composited into a watery environment. Various painted texture maps have been applied to the model to create the appearance of his skin.

LEFT: Two texture maps used to add colour and detail to Abe's body. These two were painted for his chest; many more were created to cover every part of his body.

FAR LEFT: While animatronic creatures were used for many shots on *Hellboy*, Tippett Studio created CG versions when they were required to leap or jump, such as this shot of the multi-tentacled character Sammael.



Created by Weta Digital, Gollum was one of the first digital characters for which subsurface scattering was used to produce incredibly subtle and lifelike skin.

DIGITAL SKIN

One of the greatest barriers to creating convincing computer-generated characters has been the ability to synthesize natural-looking skin. In 1993 ILM created photorealistic digital dinosaurs for *Jurassic Park*. These creatures had thick, leathery skin that was perfectly suited to the computer's ability to render solid materials from which light would bounce directly. But eight years later the beautifully modelled computer-generated humans of *Final Fantasy: The Spirits Within* (2001) remained unconvincing because light was still bouncing off their skin, making it look solid and waxy.

To solve this problem a new form of shader model has been developed that accounts for the fact that organic materials such as skin contain large amounts of water, through which light easily passes. The new technique, called 'subsurface scattering', is a mathematical description of how light that penetrates the surface of skin is absorbed or scattered by layers of subcutaneous flesh and blood before re-emerging at varying angles and with different colours (fig. 5). The result has been subtle and naturalistic computer-generated skin for characters such as ILM's Dobby the House-Elf in *Harry Potter and the Chamber of Secrets* (2002) and the green star of PDI DreamWorks' *Shrek 2* (2004). One of the first examples to make it to the screen was Gollum, the extraordinarily convincing character created by Weta Digital for *The Lord of the Rings: The Two Towers* (2002).

'When I took on Gollum I knew that creating a believable 3-D animated performance was no longer a major issue and that the key to believability was going to come down to the quality of his skin,' states Joe Letteri, who won Oscars for his work as visual effects supervisor on both *The Lord of the Rings: The Two Towers* and *The Return of the King* (2003). As he began working on Gollum, Letteri discovered the new research into subsurface scattering and thought it might provide the solution he was looking for. 'People were starting to experiment with subsurface scattering,' recalls Letteri. 'However, in order for it to work they were actually modelling everything under a character's skin for the light to interact with, including capillaries, blood vessels and fatty tissues. This seemed a tremendous amount of work for details that would never actually be seen on the screen. Then one day I saw a silicone model of Sean Bean's head that had been made by Weta Workshop. It looked so real – as if the actor's head had just been chopped off. The skin looked like it had great depth and that all those natural details were actually lying below the surface.' Letteri learned that the head was made of solid silicone that was dyed a base flesh colour. On top of the silicone a number of thin layers of colour pigment had been hand-painted to build up the impression of depth.

'I realized this might be a practical method of creating realistic skin for Gollum,' recalls Letteri. 'The method we devised as a result involved developing a base digital flesh material that was equivalent to the silicone used for the model head. Then we layered a number of detailed colour and texture maps on top to form the skin. In this way light would first pass through Gollum's skin layers picking up colour detail, then enter the flesh layer where the photons of light were absorbed or scattered before exiting back through the skin layers to pick up more detail.'

The result was the most believable photorealistic digital character yet seen on film. 'Subsurface scattering really did add something to our creative toolkit,' remarks Letteri. 'We could put Gollum in bright sunlit scenes next to Sam and Frodo and his skin would just soak up the sunlight, making him look just like his human co-stars. It's a breakthrough that adds a vital level of realism. We can now push light into skin and get real softness and richness. We used it for all our CG characters in *Return of the King* and *King Kong* [2005].'

FIGURE 5 SUBSURFACE SCATTERING

Normal rendering techniques (a) calculate the way that light bounces off the surface of objects. Subsurface scattering (b) simulates the way that light penetrates translucent materials, being absorbed and refracted before exiting with different qualities. Subsurface scattering is capable of producing incredibly realistic-looking CG skin.


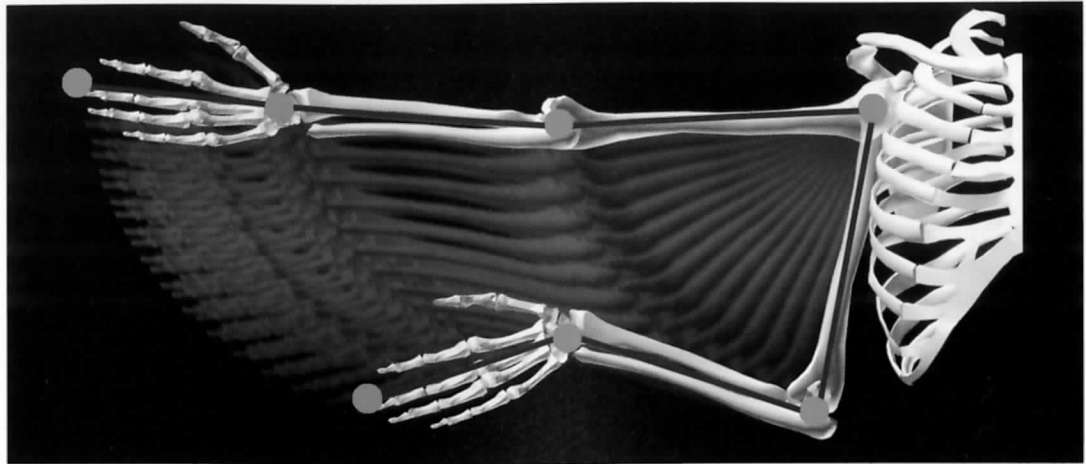


FIGURE 6 KEY-FRAME ANIMATION

As the result of setting just two key frames at either end of a large movement, the computer automatically calculates each position in between. To produce this lifting arm motion, for example, the computer animator need only move the model twice. A traditional stop-motion animator might move their model six times to produce the same effect.



ANIMATING DIGITAL CHARACTERS

A single second of film comprises 24 individual frames. A stop-motion animator must therefore painstakingly position and photograph a puppet 24 times in order to produce just one second of movement. At best, traditional animators might work 'on twos', which means that they expose two frames of film at once, thereby cutting their workload in half. But if the animators make a mistake, they may have to discard hours of work and start again.

Computer animators have several distinct advantages over traditional stop-motion animators. Since computers can mathematically interpolate movement, animators can create a minimum number of character positions per second and rely on the computer to intelligently calculate the rest for them. Furthermore, a computer animator can refine work endlessly, gradually building up a performance and only finishing a piece of animation when it has been approved by the client.

The computer animator's primary tool is the 'key frame'. Each key frame is the point in time at which an animator manually changes one aspect of the object or character that they are bringing to life. Creating animation in this way is called 'key-framing' and it is the technique used to produce most computer animation.

An animator first informs the computer how many frames there will be in a shot. This creates a timeline with each frame in the sequence marked along its length. The timeline, displayed across the bottom of the screen or on a separate monitor, is used to place key frames for all of the changes made to all of the objects in the scene. Every changeable aspect of each object will have its own track in the timeline. For example, a model dog might have separate tracks to control the motion of its body, head, eyes, ears, tongue, tail, legs, feet, and any of the other bodily functions that have been rigged for animation. The motion of every one of these elements can be addressed at any point by accessing its track in the timeline and adding, deleting or changing key frames. Key frames may also be used to control other aspects of a computer-animated scene, such as the camera and lights. A simple computer-animated sequence may ultimately have several dozen key-frame tracks while a complex scene with several characters may have hundreds of tracks, each of which can be accessed and altered at any time.

The animator begins by arranging the object in its first pose. Computer animators use an ordinary computer mouse to drag and position the joints of a digital model using handles that have been attached to the skeleton by the riggers. Some aspects of a character's motion can also be controlled by 'sliders' that animators use to dial in the performance they want. Sliders might be used to control the speed at which a character breathes, or the rate at which a dog wags its tail, for example. When the first position of the object is satisfactory, the animator sets the first key frame. The animator then moves the object to its

second position. When happy with the second pose, the animator sets the second key frame.

The animator continues through the sequence, setting key frames whenever the object needs to change position significantly. If the object being animated is moving linearly, a key frame may be set perhaps once every 12 frames. With the start and end positions of a move identified, the computer interpolates the two poses and automatically 'in-betweens' – that is, calculates each of the incremental movements that need to take place in the intervening frames to produce smooth movement (fig. 6). Objects or characters with erratic movements that alter radically from one frame to the next may be key-framed for every frame of a sequence.

Animators typically deal with a single aspect of a performance at once. If animating a bird in flight, for example, the overall motion of the bird will be created first. If flying in a straight line this movement may be defined using just two key frames – one for the starting position of the body and one for its end position. Next the animator might create the more complex motion of the flapping wings, followed by the turning of the head, opening of the beak and swivelling of the eyes. Continuing in this manner, the animator gradually builds up the detail and complexity of a performance.

The fact that computer animators are able to repeatedly review and refine the movement of a character does not mean that creating a performance within the computer is any easier than when using traditional techniques. 'Many people think that computer animation is much simpler than stop-motion,' remarks Todd Labonte, an animation supervisor at Tippett Studio, 'but it's still all about creating a performance. Animators still have the dilemma of deciding how, when and why they will create certain movements. The only thing that has changed is the tool used to translate those decisions.'

'In some ways animating a convincing performance is harder with the computer,' continues Labonte. 'For a start, using an external device like a mouse to move an on-screen character is hardly as intuitive as manipulating a stop-motion puppet with your hands. In fact, I often feel like I'm wearing boxing gloves while I'm animating!'

When animators receive a rigged character model, the first thing they determine is the way it walks. 'Walking is used as the basis for all character performances,' says Labonte. 'This is where we really define the way a character moves. For example, a lion and a domestic tabby cat have pretty much the same basic physique. But the way we animate their walk will make that considerable difference between a large powerful-looking animal and a small harmless pet. Walk cycles show the way a character holds their body, what their weight and size is, and are a glimpse into how their brain works.'

During production several animators will often work with the same character simultaneously. To ensure that a performance is consistent from one shot to the next the animation supervisor will therefore create a character



'bible', demonstrating how a character poses and moves in different situations. 'As animation supervisor I'm normally the first animator to get hold of a character,' says Labonte. 'I then work with the movie's visual effects supervisor and director to lock down the way a character performs. This will include things like walks and runs, whether the character moves first from the shoulders or the hips, any small habits like the way they twitch their hands when standing still, and so on. I then produce a number of sequences and illustrations that I share with the animation team. I'm essentially an acting coach who has to make multiple animators deliver the same performance.'

Despite such guidance, animators do have their own distinctive style and so each shot will be assigned to someone considered most suited to produce the necessary motion. 'As supervisor I have to "cast" each shot,' remarks Labonte. 'Sometimes I'll read a script and know instantly that a fast-moving piece of action should be handled by one animator, while a subtle piece of emotion should be given to another.'

Once a shot has been assigned, the supervisor and animator will discuss what is required. 'We will talk through the shot, look at the storyboard, discuss what the director has requested and throw any new ideas around,' explains Labonte. 'The animator will then make a start by blocking out the basic motion of a character. Then we'll show that to the director for approval before starting on a refined performance. If the character is going to be composited into a live-action shot we will need to see how it will fit into the pre-filmed background plate. Often a director will ask for a particular action but then we'll find that it doesn't really work with the camera movement in the live-action plate. This is usually because some poor camera operator has had to guess how an as yet invisible character might look when moving through the frame. Then we'll have to design a performance that does what the director wants and also physically fits into the shot.'

In order to create a performance, animators will often act out their character's motions and videotape them for reference. They will also look at footage of real animals in action. 'Most animators have a mirror near their workstation so they can look at themselves performing,' says Labonte. 'It's important that an animator can feel what their character is supposed to be doing and can identify where an action or emotion is coming from. I'm sure most animators are frustrated actors who have some kind of hang-up about putting themselves in front of the camera. If you look around the studio you'll see the animators at their desks repeatedly, almost subconsciously, acting out little movements and expressions. I've gone home with a sore neck many times!'

Ultimately a performance is revised numerous times as the result of suggestions made by many people involved in the production. 'We review shots each morning in dailies,' explains Labonte. 'We'll loop a shot and watch it a hundred times so that everyone can study the work and make suggestions. I then have to filter those ideas and decide what changes to make before finals are sent for approval to the director.'

During the animation of a character, animators can view their work in a number of ways. These images from the production of *Evolution* (2001) show how a character was animated by Tippett Studio.

TOP LEFT: The basic computer-animation user interface. The character can be seen from four 'orthographic' views. The timeline for the scene runs across the bottom of the screen, and at the top a toolbar contains various animation tools.

TOP RIGHT: To check that their animation works within the live-action scene, the animator can superimpose the character over the live-action plate using the correct camera perspective.

ABOVE LEFT: The animator can give the character a simple shaded surface, concealing the wire-frame model and giving a clearer image of the performance.

ABOVE RIGHT: The final fully rendered shot as seen in the film.

PERFORMANCE CAPTURE

Since the earliest days of motion picture photography and animation, artists have sought ways to capture and reproduce naturalistic character motion. Eadweard Muybridge (<11) developed a system that enabled him to photograph and study the momentary poses of moving creatures; Max Fleischer designed the rotoscope (<174) so that animators could copy human motion for use in cartoons. Today, most computer animation is achieved using a mouse to manipulate digital models and set key frames within the digital environment. However, the idea of harnessing real movement is still attractive to film-makers, and a number of methods of channelling externally driven movement into the computer have been developed.

Initial plans for creating the dinosaurs in *Jurassic Park* (1993) involved a combination of full-scale robotic dinosaurs and traditional stop-motion puppets animated by Phil Tippett. Tippett's hand-animated dinosaurs were to be digitally composited into real landscapes by Industrial Light and Magic, who planned to give them added realism by applying digital motion blur (239>) to their movements.

Although at the outset of production digital technology was not considered capable of creating convincing dinosaurs, ILM's senior visual effects supervisor, Dennis Muren (<47), was convinced it would be possible to produce computer-generated dinosaurs of the quality required. Muren produced an early test that convinced Steven Spielberg of the potential of the computer. As a result, the task of creating 50 dinosaur shots was moved away from the traditional method of stop-motion and into the relatively unexplored territory of the computer.

Tippett was devastated that the computer had made its first serious incursion into the world of the stop-motion animator. Like many other artists, Tippett feared that his skills would soon become redundant. However, while ILM's computer wizardry was able to create convincing dinosaurs, the technology was not yet capable of producing reliably good performances. 'The first test I produced for Steven [Spielberg] was of a running tyrannosaurus,' recalls Muren. 'Steven was bowled over by the first test and gave us the go-

ahead to produce all of the movie's dinosaurs in CG. But after our initial test we had real problems getting our CG dinosaurs to perform well. We started to worry that our test had been a fluke and we wouldn't be able to pull off a whole movie full of CG dinosaurs, so we spoke to Phil Tippett about ways to fuse traditional animation techniques with our new technology.

'We realized that there had to be a way to tap the experience and art of stop-motion animators and feed that into the computer,' says Tippett. 'By combining the two disciplines, we thought pretty amazing things ought to be possible.' In conjunction with ILM, Tippett and his staff built an interface between computer and traditional animator that they called the DID (dinosaur input device, later refined to direct input device).

'The DID was basically a traditional stop-motion armature,' states Tippett, 'but instead of being covered with foam latex and painted to look like a dinosaur, it remained just a skeleton. Each joint of the skeleton had a number of encoders that could measure the degree and direction of movement at that pivot point.' The DID device was linked to the computer, where a wire-frame model of the dinosaur had been built with corresponding joints. Using the DID, stop-motion artists animated a dinosaur's performance in the traditional way, then saw that movement translated to the digital model within the computer.

The process combined the best of both traditional and cutting-edge techniques. Animators were able to manipulate the armature to produce a nuanced performance, and unlike the old method of capturing movement on film, the new method allowed the animator to go back and refine any moves which they were unhappy with. After basic animation was approved, ILM added the digital muscle and skin to create lifelike animals. The DID device was used to create the terrifying tyrannosaurus attack on the *Jurassic Park* jeeps, and the sequence in which two children are stalked by velociraptors in the kitchen.

Following *Jurassic Park*, Tippett Studio animators used DID technology to help produce the stunning insect performances in *Starship Troopers* (1997). However, with the development of more sophisticated and intuitive software, today's animators now create subtle character performances using only their mouse to manipulate on-screen skeletons. Devices such as the DID are no longer necessary.

RIGHT: Craig Hayes from Tippett Studio with the Dinosaur Input Device (DID) that he helped design and build for the production of the pioneering CG dinosaurs of *Jurassic Park* (1993).




Although the motion of computer-generated creatures is generally created by animators, the complex and naturalistic motion of humans is normally captured directly from real life. In the 1980s, various medical establishments created methods of digitally recording and analysing the movement of patients who suffered from leg and spinal problems. This technology, known as motion capture, or 'mo-cap', has since been harnessed and developed for use in visual effects production.

'There are two main methods of capturing human movement for use in computer animation,' explains Tom Tolles, CEO of House of Moves, a company that specializes in capturing movement data for use by visual effects facilities. 'Optical motion capture uses a number of cameras to "film" movement and translate it into computer data. Magnetic motion capture, on the other hand, uses a strong magnetic field to measure the movements of a person dressed in a bodysuit that is covered with magnetically active markers. Magnetic motion capture is generally more restrictive and complicated to use, but it does produce excellent data that can be utilized to create real-time performances. This means that a performer's actions can instantly affect the movement of a computer character in live-broadcast applications such as kids' TV, where animated characters need to interact with an audience.'

House of Moves uses the more popular optical method to capture movement. 'Our system works by using somewhere between 4 and 250 4-megapixel CCD [charge-coupled device] video cameras arranged around a performance area that is perhaps 8 m [26 ft] in diameter,' states Tolles (see fig. 7 for basic set-up). 'The performer is dressed in a black suit to which we attach a number of small plastic balls. The balls are covered in a highly reflective material similar to Scotchlite, which was once used for front-projection purposes [<84]. Depending on the sophistication of motion detail needed, performers are dressed with between 20 and 60 markers at all the key joints. For basic movement, a hand can be represented by two markers, one at the wrist and one at the tip of the hand. Really complex movement can involve having a separate marker to capture the movement of each joint of each finger. Optical motion capture can also be used to capture the subtle movements of a human face. Up to 300 markers can be attached to the skin and lips to capture the expressions of a performer.'

FIGURE 7 MOTION CAPTURE



During the capture process, performers are filmed moving within the capture area. If their character needs to interact with any objects or surfaces these will be built in the studio – sets of stairs for them to walk up or fake cars for them to step out of, for example. Each CCD video camera emits infrared light, which hits the reflective body markers and is directed back into the camera it came from. The result from each camera is a black image containing a number of moving white dots – a 2-D representation of movement as seen from a single position. Seeing this cloud of dismembered floating points is a strange experience. Although nothing connects the markers, their movement fools the mind into making the logical connection between dots to form the impression of human movement – even the feeling that you can see limbs that do not exist.


To extract useful information, the computer to which this data is fed compares the centre of each white dot from several camera angles and triangulates the data to calculate its exact position in 3-D space. 'The computer needs to cross-reference continually between a number of camera angles,' states Tolles. 'Looking at a performer from just one or two angles isn't enough, because each moving marker is continually being hidden, or "occluded", behind various parts of the body.'

Although the system is designed to calculate the position of each marker automatically, the resulting data normally has to be cleaned up manually. 'Motion-capture data has to be organized before it can be used,' Tolles explains. 'The markers are continually moving about. Sometimes one can pass another, and then the computer can get confused and swap them. The result could be an arm movement getting mixed with a hip movement, which would result in a pretty weird-looking performance!'

'This system produces a huge amount of data,' notes Tolles. 'Each CCD captures images at least 60 times a second, and with 50 markers on a body, that's 3,000 measurements per second per camera. However, it is because the system is so sensitive that it can capture recognizable performances. We do a lot of capture for the games industry, which hires famous sports stars to come and have their basketball or boxing movements captured for computer games characters. We've captured the distinctively recognizable moves of Bruce Willis for the action game *Apocalypse*, and of tennis player Andre Agassi for a Nike commercial. But we're best known for our work for the movies. An early mo-cap breakthrough was the work we produced for *Titanic* [1997], which allowed Digital Domain's CG characters to walk around on the decks of the model ship. More recently we've provided facial capture services for *Spider-Man* [2002] and *Spider-Man 2* [2004], and services and software for *The Polar Express* [2004] and *Monster House* [2006].'

Once motion data has been captured it needs to be applied to computer-generated characters in order to make them move. Each marker on the body of the live performer is mapped onto the body or skeleton of the digital character. In most cases the motion-capture data will need to be distorted somewhat – the markers used for motion capture are attached to the outside of the performer's body while the bones and joints that they drive in a digital character are a little way below the surface of the body. Similarly, any digital character that is a different shape or scale to the mo-cap performer will need to have its motion data mapped accordingly. For *The Polar Express* Sony Pictures Imageworks used data captured at House of Moves to drive the movement of the film's animated characters. Actor Tom Hanks played a number of characters in the film, including a seven-year-old boy. To create realistic movement, Hanks's motion-captured movements were scaled down to fit the smaller and differently proportioned body of the young boy. The props and scenery he interacted with during these scenes were also made larger to make his actions seem more like those of a young boy.

Mo-cap data is often only the starting point in the production of a performance. Once the natural movements of a character have been captured they can be altered using key-frame animation to produce a more dramatic or subtle performance. A character's motion is often captured in a series of discrete movements to build up a 'library' of moves. These small moves can later be edited together to produce a synthetic performance that is a combination of a number of natural movements. This technique is particularly



useful when creating large crowd scenes in which many characters need to move in reaction to various influences. In such cases the computer can judge what motion a character needs to make and assemble the appropriate pieces of motion-capture data to automatically create a performance (216>).

Motion capture has been used with great success in certain situations, but it does have its limitations. 'Mo-cap is really good at getting naturalistic human performances into the computer,' believes Tom Tolles, 'but some people tend to think that it is an easy way to achieve motion for all sorts of different animated characters. If the motion of a real animal is needed, such as that of a horse or a dog, we can bring those animals into the studio and record their movement just like that of humans. If it's a fantasy character being played by a human we can load performers with heavy weights so that they move in a slow, ponderous way – more like a giant or a monster. However, fantasy creatures are probably best created using key-frame animation techniques, which gives them a stylized performance. The monsters of Ray Harryhausen probably look so fantastic and appealing precisely because they have such a mysterious and undefinably inhuman way of moving.'

ABOVE LEFT: A stunt performer covered with tracking markers slides down a tilted deck, providing motion-capture data for the creation of a computer-generated passenger in *Titanic* (1997).

ABOVE: Even the *Titanic's* child passengers had their movements captured. This girl is wearing tracking markers on her skirt: the swaying motion of the skirt will be transferred to the digital clothes of her animated double.

LEFT: For this shot from *The Polar Express* (2004), Tom Hanks played all three characters. His movements were recorded in separate motion-captured performances applied to the different characters and later combined in a single shot.



JURASSIC PARK

Early plans for Steven Spielberg's (<39) adaptation of Michael Crichton's novel *Jurassic Park* involved a combination of full-scale animatronic dinosaurs and traditional stop-motion techniques. However, in the wake of its success with digital characters in *Terminator 2* (1991), Industrial Light and Magic was able to convince Spielberg that digital dinosaurs were a viable alternative to stop-motion.

To create their digital dinosaurs, the artists at ILM, under the guidance of effects guru Dennis Muren (<47), employed a number of groundbreaking techniques. The dinosaurs were first created as clay sculptures and then cyberscanned (<195) to produce digital models. These were then given a digital skeleton so that they could be animated using inverse kinematics (<201). ILM programmers created new 'Enveloper' software, which allowed the skin of the models to wrap round and move freely over internal muscles in order to create the look of a real animal.

ILM also created a program called 'Viewpaint', which allowed artists to paint the surface texture details of a dinosaur's skin directly onto the digital model. Such texture maps (<164) would previously have been painted as a flat design and then wrapped around the digital model, often causing stretching and distortion of the design. Viewpaint allowed skin designs to be painted in their correct place on the model itself, much like painting a physical model with paints. Several layers of texture map were applied to each dinosaur model, including bump maps to describe the texture of the skin and additional maps to describe skin colour and pattern, mud and dirt. For the sequence in which the T rex attacks jeeps in the rain, a special shader was written to create the effect of water running down the creature's back.

Computer animators created the dinosaur performances by referring to simple animatics – low-resolution test shots that were created and refined using traditional stop-motion puppets or basic computer models. The final dinosaurs were then animated using a combination of key-frame techniques (<205) and a 'dinosaur input device' (<207), which allowed stop-motion animators to transfer their traditionally animated performances directly into the computer.

Once dinosaurs had been animated and rendered (236>), they were composited (<100) into live-action background plates. The compositing process was vital to the credibility of the dinosaurs – even the most realistic digital characters will look fake if they do not convincingly integrate with their surrounding environment. As the T rex ran, small clouds of dust and dirt were added beneath its feet, shadows were cast on the ground and splashes of water were placed where it stepped in puddles. For shots in which dinosaurs emerged from behind objects each frame of footage of live action was rotoscoped by hand (<174). Individual leaves and blades of grass were painstakingly traced so that the digital creatures could be sandwiched into natural environments.

As production progressed, ILM became more confident in its ability to deliver photorealistic dinosaurs and the creatures were given a greater on-screen presence. A late addition to the schedule was a scene in which the actor Martin Ferrero is eaten by the T rex. According to the original plan the character would be seen disappearing out of frame when snatched from his seat on a toilet. The revised version showed the actor being picked up, tossed around and then eaten by the dinosaur in a single shot. During compositing, the real performer was swapped for a digital character just as the dinosaur's mouth closed over his body. Though a last-minute decision, the shot became a show-stopper.

During pre-production planning, digital dinosaurs had only been intended for use in long shots, with Stan Winston's animatronic puppets used for close-ups. However, as ILM's work became increasingly convincing, the CG creatures were brought closer to camera. For the last sequence featuring the T rex, the dinosaur's texture maps were repainted with extra detail and the creature brought right in front of the camera as it rampaged through the Jurassic Park visitor centre.

The public flocked to see *Jurassic Park*, smashing box office records to make it the highest-grossing film in history. The film later won Academy Awards for its visual effects and sound effects. Perhaps the film's greatest legacy was to prove that, finally, the computer could help special effects artists create almost anything the scriptwriter could conceive. Film-making was changed forever.





FACIAL ANIMATION

Expressive facial movement is the key to any convincing character performance and is among the hardest forms of animation to achieve successfully. Most computer-generated facial animation is currently produced using one of two different techniques. One method creates a physically based digital model that emulates the way muscles and flesh move on a real face. The other technique models the surface of the face as a series of fixed expressions. These are then arranged in a sequence to create a performance – the digital equivalent of replacement animation (<191).

Physically based animation requires the building of complex facial rigs that resemble the mechanics of the human face. Beneath layers of digital skin, tissue and ligaments lies a network of up to 200 muscles. Each muscle has a controller which allows it to be pushed and pulled by the animator. This movement translates through the layers of tissue to deform the high-resolution facial geometry of the face surface. The animation of each muscle usually affects the shape of the surrounding muscles so that if the mouth is moved, for example, the cheek will deform in response.

Individually animating each of the muscles used in such faces would be an arduous task and so interrelated muscles are normally clustered into a series of control groupings. Using these, animators can move discrete parts of the face to produce key facial movements such as 'right eye close', 'lip sneer', 'forehead wrinkle' and so on. These movements are key-frame animated to create changing facial expressions throughout a sequence. Techniques like these were used by Pixar to create facial animation for films such as *Monsters, Inc.* (2001) and *The Incredibles* (2004).

The animation of physically based face models can also be driven by motion-capture data derived from the performance of an actor. This technique was used to create the facial performances in *The Polar Express* (2004) and *King Kong* (2005). Commonly called 'performance capture' rather than 'motion capture', this method requires an actor to wear dozens, sometimes hundreds, of tiny facial markers. The movement of these markers during a performance is recorded like normal motion capture (<208) and then mapped onto the muscle groups in a digital facial rig in order to create an animated version of the original performance. Motion capture normally only provides perhaps 80 per cent of a performance, the remaining animation being created with normal key-framing techniques. In cases where the live performer's features do not directly translate to the geometry of the digital character's face, the motion-capture data will need to be distorted somewhat. For *King Kong*, for example, the motion around performer Andy Serkis's nose and mouth needed to be remapped to better fit the wide, protruding muzzle of the great ape.

The other popular method of producing facial animation relies on creating a series of pre-sculpted facial expressions. With this method the face model has no underlying flesh or muscle and so the geometry of its surface is directly sculpted using traditional digital modelling techniques. Sometimes each expression is first sculpted in clay and then scanned. Animators then select sculptures showing the expressions they want to use and place them at the key frames at which they are needed. The

ABOVE LEFT: Tom Hanks undergoes performance capture for a scene from *The Polar Express* (2004). Dozens of tiny markers capture the movement of his face, while larger markers on his head, hands and shoulders are used to relay his overall body movement.

ABOVE: The finished scene on film.

RIGHT: Animators have traditionally used their own reflections in order to help them create convincing facial expressions for their characters, but to help his animators produce the right performance for Yoda, ILM animation director Rob Coleman created video reference shots of himself. *Left*: Here Coleman acts the part of Yoda for a shot in *Revenge of the Sith* (2005). *Centre*: The computer animator combines a number of blend shapes to give the CG Yoda a similar expression. *Right*: The shot as seen in the final film.

computer then interpolates or 'blends' from one expression to the next over a designated number of frames to produce the required facial movement. This technique, known as 'blend-shape animation', was used to produce the facial performance of Gollum in *The Lord of the Rings* and Yoda in *Star Wars: Episodes II and III* (2002, 2005).

'The face of our CG Yoda began as a single digital sculpture in a neutral pose,' explains ILM animation director Rob Coleman, who oversaw character animation on all three modern *Star Wars* movies. 'From that starting point we digitally resculpted segments of Yoda's face to produce 40 separate facial elements, or "blend shapes". For example, two of our 40 blend shapes were a "right eyebrow up" shape and a "right eyebrow down" shape.'

Each of the 40 blend shapes was assigned a digital slider, creating an on-screen control panel similar to that of an audio mixing desk. Each blend shape in the animator's timeline could be faded up or down between 0 and 100 per cent. The higher the number, the more that shape influenced the overall facial pose. 'Using these sliders we were able to blend between our 40 differently sculpted shapes to produce Yoda's entire facial performance, including his speech,' says Coleman.

Creating a performance of nuanced emotions, expressions and speech from just 40 sculpted elements would seem like a tall order given that the human face uses hundreds of muscles to produce its range of movement. Yet ILM has found that it is not the number of shapes used but the way they can be selectively combined that creates a naturalistic performance, as Coleman explains: 'For earlier films we used to sculpt entire facial expressions. If we wanted an angry face we would have a bunch of models in which the whole face was pre-sculpted with different degrees of anger. There might be hundreds of fully sculpted faces altogether. During animation, one complete facial expression was blended into another. Performances created in this way worked well for non-human characters but they tended to be rather stylized and didn't really have the range and spontaneity that you find in a real face. We now have a much better understanding of how facial movement works, and by intelligently designing and then blending a relatively small number of facial elements, rather than overall facial sculpts, we could produce a very organic and expressive performance for Yoda. In actual fact, of the 40 blend shapes made for Yoda we used only 25 of them continually while the other 15 were used only occasionally.'

Although Yoda's facial expressions were created by a selective combination of blended facial elements rather than sculpted full facial poses, Coleman's team did rely on a number of pre-designed facial models. 'Our lead animator, Jamy Wheless, blended shapes to create eight master expressions,' says Coleman. 'Any time one of the fifteen animators working on Yoda wanted to create a particular emotion they could just press a button and up would pop the "meditative" face or the "angry" face. All the sliders would move to the correct blend positions to create these expressions. From there the animator could start blending shapes to create a performance, but it meant that everyone had a common starting point and could stay "on-model".'

Coleman describes how a Yoda facial performance was typically created: 'The first thing we do is block in the head action. This means simply key-framing the changing position of the head throughout the shot according to where he needs to look. We then listen to the recording of Yoda's dialogue and animate the jaw to indicate when he is or is not talking. This is simply jaw down, jaw up. The jaw controller is not one of our blend shapes, it's just a normally animatable part of the face.'

The next stage is to start layering in more specific mouth movement to produce the appearance of a voice performance. 'First we put in the mouth-closed phonemes – the Ms, Bs and Ps,' says Coleman. 'We then work on the mouth-open positions. At this stage we don't put in the individual phonemes, we start with the loudest sounds in the shot – where the mouth is at its widest. With those layered in we have an overall shape to the performance – we know when the mouth is open and closed. We then create the finer lip movements for the remaining phonemes – the Os, Es, Us and so on, in order to synchronize exactly with the pre-recorded dialogue. This is one of the tasks that uses the most blend shapes. For example an E will require three blend shapes to influence the form of the upper lips to a greater or lesser extent, then three more for the lower lips, plus we'll have the cheeks going up on either side and also a little crinkling at the corners of the eyes. With the basic words formed on Yoda's lips we would then add in any additional head moves that might be made in response to the dialogue. English speakers tend to emphasize certain words with a slight nod or tip back of the head.'

Finally the animators must give some thought to the performance of the upper portions of the face. 'Animating the mouth is just about getting the mechanics of talking right,' says Coleman. 'But the eyes are what really convey emotion. This is partly animating the face around the eyes – the way the brow furrows or the eyelids squint – but also about how the eyes themselves move.'

A character's eyes can be animated by using a controller that works like a handle protruding from the iris of each eye. The end of the handle can be moved around to sweep over the places that the character needs to look. The handles can also be attached to whatever the character is looking at so that their gaze remains fixed even as they walk around.

'Eyes are rarely still, even when we stare at a person's face during a conversation,' remarks Coleman. 'If the eyes don't move a character looks lifeless so we always have them scanning back and forth a little. In addition the pupils can widen or dilate to convey thought and emotion. We can also alter the convergence of the pupils, making the gaze of the eyes parallel when a character stares into the distance or cross slightly when staring at something very near.'

The hardest form of animation to create convincingly is undoubtedly that of the human face. We are all experts in what faces of all shapes, shades and emotions look like and are not easy to fool with any form of synthesized performance.

A startlingly realistic form of human facial animation was developed for use in *The Matrix Reloaded* (2003). To produce scenes in which the film's stars appeared to fight and perform in ways impossible for a real human,



a team led by George Borshukov was asked to find a way to create photorealistic computer-generated doubles that could be seen in close-up.


The technique they developed, called 'Universal Capture', was an image-based animation system. This meant that it used images of real moving faces to drive the synthetic motion of a computer-generated face. The CG face could then be attached to an animated body and be viewed from any angle under variable lighting conditions.

The system worked by surrounding performers with an array of five synchronized high-definition digital video cameras (fig. 8 (a)). Actors were then filmed as they performed the necessary expressions or dialogue. The result was a series of shots of the same action as seen from five different angles (b).

In the computer a high-resolution cyberscan (<195) of the performer's face (c) was then surrounded by five 'cameras', each of which was aligned to look at the digital head from exactly the same position as they had viewed the real head during filming (d). Each camera contained a record of the live-action footage that it had filmed from that same position during filming.

Each vertex on the surface of the cyberscanned face model was then 'projected' into each of the five surrounding cameras. Within each camera, each vertex from the surface of the 3-D cyberscan lined up with the 2-D coloured pixel that had been recorded at that exact spot on the face of the performer during filming (e). As the performer's face moved from one frame to the next, the changing position of every coloured pixel, and the vertex on the face of the cyberscan which it represented, was tracked. By triangulating the movement of each pixel from the position of each of the five cameras, an accurate reconstruction of the path of each pixel, and therefore vertex, through 3-D space was recorded. The motion of the vertices was then used to deform the surface of the cyberscanned head in order to create an animated performance that exactly copied that of the actor (f).

To complete the animation the 2-D video images filmed from each of the five camera angles were combined to produce an animated texture map that fitted the moving facial model perfectly. The result was highly controllable digital doubles that were indistinguishable from their real-world counterparts.



LEFT: The eight master expressions of a Jedi Master, as used by ILM when animating their digital Yoda.

BELOW: Universal capture allowed dozens of photorealistic computer-generated clones of actor Hugo Weaving to appear in this scene from *The Matrix Reloaded* (2003).

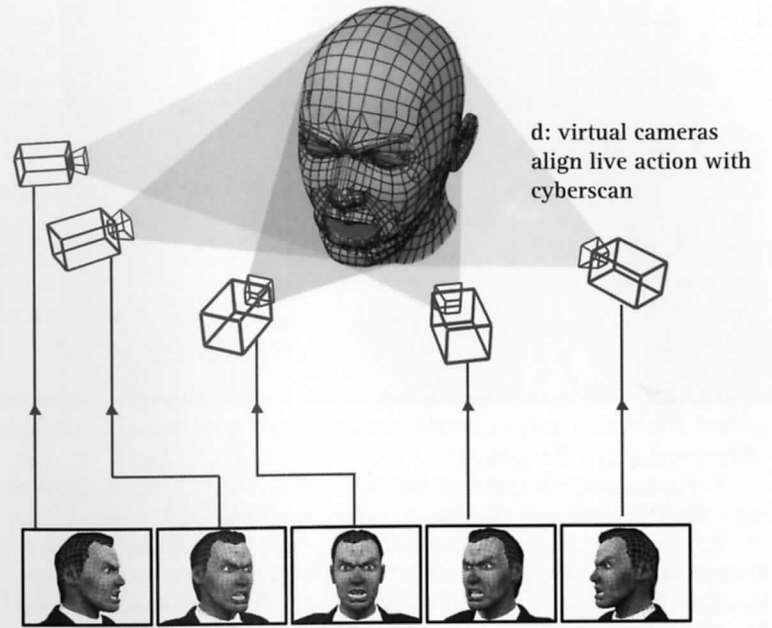


a: performer filmed by high-definition cameras




b: performance captured from five different angles

c: cyberscan of performer's head



d: virtual cameras align live action with cyberscan

e: live-action performance used to drive animation of cyberscan



f: computer-generated character replicates actor's performance.



CROWDS

Film-makers have long realized that one of the most exciting movie spectacles can be the sight of hundreds or thousands of people involved in the on-screen drama. In 1915 D.W. Griffith employed 5,000 extras to populate the Babylon sequences of *Intolerance*, while an impressive 20,000 troops were recruited from the Soviet Red Army to re-enact Napoleonic battle scenes for the epic flop *Waterloo* (1970). But the record is held by Richard Attenborough's *Gandhi* (1982) which rallied over 250,000 extras for its funeral scenes.

But filling the screen with legions of people is time-consuming and very expensive. Because of this, film-makers have sought ways of making relatively few extras look like many more. Digital split-screen techniques now allow groups of extras to be filmed in several positions and then combined to look like a large crowd (<108). Even in the age of digital deception there are still some brilliantly simple methods of creating a crowd. For *Seabiscuit* (2003) and *Wimbledon* (2004), film-makers used hundreds of inflatable dolls wearing masks, clothes and wigs. Providing the camera was moving or the crowds were in the background of a shot, the conceit was undetectable.

However, some productions have called for tens and even hundreds of thousands of characters to perform in incredibly complex battle scenes. To produce these, film-makers have turned to the image-generating powers of the computer.

The largest movie battles ever created appeared in Peter Jackson's *Lord of the Rings* trilogy, with some sequences featuring over 200,000 battling orcs and elves. To produce these scenes Weta Digital created a unique crowd simulation software program called 'Massive'. 'Massive combined digital character animation techniques with a form of artificial intelligence that allowed us to control how thousands of characters looked and behaved,' states visual effects supervisor Joe Letteri. 'We built a kind of artificial brain which was a system of rules governing how characters might behave or react in various situations. The brains were a network of 7,000 to 8,000 nodes, each node being the equivalent of a decision that had to be made, such as "Do I lift my sword or not?" or "Is this person so strong that I should run away?" Our characters could recognize who to attack, which weapons to use, and how to stand, run, or fall on different types of surface. Warriors even knew how to die in a manner suited to the way they were attacked and where they were standing.'

Once programmed, characters were let loose in environments that contained various obstacles and enemies. The computer then simulated a battle with each character behaving according to its own rules and drawing on a library of discrete motion-captured body moves to generate their actual performances (217>). 'Because we only told agents how to react in certain situations rather than exactly how to perform in every frame, the resulting battles were quite organic and could contain all sorts of surprises,' says Letteri. 'In one shot a group of characters were so smart they actually decided they didn't want to fight and they turned around and ran away!'

Other notable computer-generated battles have appeared in *Troy* (2004) and *Kingdom of Heaven* (2005). In each case the fighting hordes were created by London's Moving Picture Company (MPC). 'For each of these films we had to produce dozens of battle shots in a very short time,' states MPC technical director Carsten Kolve. 'Most shots needed to convey a specific part of a battle rather than just show a wide vista of thousands of combatants randomly fighting. Because we essentially needed to art-direct each shot, the artificial life solutions used for the *Lord of the Rings* battles weren't really appropriate for us. We couldn't spend the time programming all the variations and running simulations – we needed to determine exactly where and when events would occur. To solve this we created our own crowd simulation engine.'


MPC began by plotting every conceivable action that might be required on the battlefield, from troops waiting for battle to begin, to charging, hand-to-hand combat, throwing spears, climbing ladders, and so on.



Historical epics often call for a cast of thousands. Whereas once many hundreds of extras had to be hired and clothed, modern crowds are more often created by the computer, as in these shots from *Kingdom of Heaven* (2005, above) and *Troy* (2004, below), in which the majority of the soldiers are digital.

ABOVE RIGHT & FAR RIGHT: For these shots from *Kingdom of Heaven*, the Moving Picture Company used a library of motion-capture moves to make thousands of computer-generated soldiers fight and die according to specific pre-programmed rules.





'Once we had a list of 500 or 600 actions we motion-captured stunt performers as they enacted them for us,' says Kolve. 'Where the action involved a prop, such as a swinging sword, we also captured the motion of the prop. In the end we had a library of many hundreds of discrete movements. We then grouped these movements together into motion libraries. A motion library is a set of movements that a character, which we call an "agent", can perform, like running to a halt, lifting and throwing a spear, and so on. The actions within a library are connected to each other with blends, allowing the movements to merge subtly from one to another over several frames. A motion library would typically contain only a small number of actions. Rather than enabling our agents to do absolutely anything in a shot, like the agents in *Lord of the Rings* – which is time-consuming and complex to control – we made our agents very good at performing specific tasks.'

In order to maintain the visual complexity of a crowd, MPC's artists created hundreds of different motion libraries that were distributed among the agents in a scene. Since shots are rarely longer than a few seconds, each agent's possible range of motion did not need to be very broad.

'For *Kingdom of Heaven*, Ridley Scott wanted very specific types of fighting to be happening at particular times in each area of a battle,' says Kolve. 'For example, at a certain point he would want the sword fighters to be battling with each other while archers shot their arrows at the people on castle battlements.' In order to arrange battles in the way the director wanted, customized fight sequences were first created by editing together the necessary sets of motion. 'We would simulate several hundred digital actors doing the right type of fighting,' says Kolve. 'We then copied or "instanced" that group of people and could rubber-stamp them across the landscape saying, "We'll have 500 here and 500 there," and so on. Each time we instanced a group the agents would automatically adapt to the new terrain. We would also rotate each new patch of people and offset their timing so they didn't look the same as their neighbours. Where groups overlapped we would go in and weed out any soldiers that merged or collided with one another, though at a distance it didn't really matter if one soldier's arm went through another's body – it wouldn't be seen.' MPC also created a library of characters performing exaggerated actions such as dying in a spectacular fashion. These were hand-placed wherever it was felt the action was a little sparse or lacking in interest.

As well as characters whose actions were determined by MPC's technical directors, each scene contained a number of characters whose movements were determined by the computer. 'We scattered a number of artificial intelligence characters in each shot,' explains Kolve. 'These characters could pick up on what the soldiers around them were doing and copy their general actions. This meant that their performance was a bit more random and always a little behind that of those around them.'

They could run around with more freedom but were programmed not to bump into or walk through other characters. It just helped each shot look more natural.'

Artificially determined motion was also used when characters needed to react directly to the physical environment around them. 'We relied on motion capture whenever we could,' says Kolve, 'but there were times when we needed movement that was far more specific. For instance, if a character fell off a ladder we would use motion-captured movement for the fall, but when they hit the ground or landed on something like a piece of siege machinery, the characters would switch over to models that were programmed to react procedurally according to the laws of physics.' These characters were driven by software that gave them the physical properties of real humans, allowing them to react naturalistically to any forces applied to them. 'When we dropped a wounded character from a collapsing siege tower he would collide with the debris and crumple and deform just like a real human without us having to animate him in any way,' says Kolve.

Each battle contained a number of different armies and regiments, each with its own distinctive dress, style of armour and range of weapons. 'We had a series of controls that influenced the quantities of each type of costume, armour and weaponry there were in a shot,' explains Kolve. 'The director could just say "more wooden shields and fewer helmets" and we would adjust the look of the scene accordingly. On top of that the computer would add additional randomness to things like the patterns on shields, how people wore their clothes and so on. You would be very unlikely to see two identical characters.'

The combination of computer automation and manual intervention was also necessary when it came to killing performers that were struck by arrows. Making arrows shot by archers find and hit a target would have been a very complex task, as Kolve explains: 'Each archer would have to pick a target, predict its motion and the flight path of an arrow and then shoot in the hope of hitting. This might be what happened in real life but it is not very controllable in terms of computer simulation. The problem is further complicated by the fact that once a person is hit it might take the computer several frames to blend the "I've been hit by an arrow" motion clip into the performance, making their reaction appear too slow. So we would manually pick a number of people and tell them where and when their death would happen. During the simulation process a "marked" agent would tell the archer of his choice "Please kill me in 32 frames and, for your information, my right shoulder will be at exactly this position at that time." The archer then only had to make sure its arrow reached the right place at the right time while the suicidal agent made sure it blended into the right death motion.'

The resulting battles were made more realistic by the addition of minute puffs of dust that were spawned each time a character's foot hit the dusty ground.

CLOTH

Creating clothes for digital characters presents a number of challenges. Garments need to fit around characters and hang from their bodies naturalistically. The material needs to move freely and yet not intersect either with the body on which it is worn, or with itself as it folds and creases. It also has to react to external forces such as wind or objects with which it comes into contact.

The most widely used form of digital cloth is made from a polygonal mesh in which the segments that connect the vertices (<154) are programmed to act like tiny springs. The degree and direction in which these springs stretch and flex can be adjusted to produce materials ranging from the lightest silk to the thickest leather. The material is programmed to act according to the laws of gravity so that it will drape over any object that it is placed on. To do this it must also be capable of 'collision detection' – it needs to detect when it is touching another object so that it will lie on that surface, or be pushed by it, rather than pass through it. So that it can be creased and folded without intersecting with itself, cloth also needs to 'self-repel'. This is achieved by making the vertices of the cloth push away from each other like repelling magnetic poles.

The final look of cloth is achieved through a combination of texture maps and shaders (<164-7). A subtle bump or displacement map makes the weave look like anything from fine cotton to heavy canvas.

Using this material, clothes can be made in several ways. An item like a shirt can be sculpted as a single object much like any other piece of geometry, or the individual pieces of cloth can be cut out according to a pattern and 'stitched' together exactly as if a real shirt were being made by a tailor.

Once clothes are made they are draped over a character and left to fall naturally. If left alone the clothes will perform procedurally, moving and creasing as the character to which they are attached is animated. However, it is rare for digital clothes to move exactly as required without any additional input from artists.

'We spend lots of time creating a performance from digital clothing,' states Juan-Luis Sanchez, a digital clothing supervisor at ILM. 'When you watch a scene from a movie like *Revenge of the Sith* [2005] being filmed you notice how the actors are constantly being preened by wardrobe people. Their capes are straightened over their shoulders or pinned back with safety pins so that they look their best in every shot. We effectively do the same thing with our digital clothing. Once the clothes are on, we pin parts of them to the body so they won't move too far. We also specify areas where we want the cloth to be more flexible or stiff, such as seams and pockets. We can then adjust various parameters, making clothes react more or less to gravity. We can add a little wind to make them billow and then thicken the atmosphere so that they flow more slowly. This all allows us to sculpt the way clothing looks and moves in order to aid the drama. For *Attack of the Clones* [2002], George Lucas asked for Yoda's digital cape to move more "romantically". This meant choreographing its movement so that it acted slightly less realistically and more heroically. Instead of moving at the same time as Yoda we often made the cape move a beat or two later. We also made it glide and swirl to the floor after a fight rather than just dropping, for example.'





PROCEDURAL ANIMATION

Like digital model-making, a number of animation tasks can be achieved procedurally (<161). Such procedural tasks are normally used to produce complex forms of motion and interaction that would be too time-consuming or complicated to animate by hand. In such cases the necessary movement is determined by a series of mathematical formulas. These algorithms, often based on academic research, are written to synthesize the physical effects of natural forces such as velocity, gravity and the dynamic flow of fluids.

Procedural animation systems allow the animator to define the material properties and behaviour of the object to be animated. They then define the variable forces that are to be applied to that object. The computer then runs a simulation to calculate the resulting animation. Several simulations are typically needed, each time altering the various parameters until the required motion is achieved.

An example of a procedural animation might be the simulation of wind blowing through the branches of a computer-generated tree. Hand-animating thousands of branches and leaves would be highly impractical. However, by describing how flexible the branches are and the way in which the leaves move in a breeze, an artificial wind force could be applied to the tree to create a naturalistic animation. By varying the force of the wind or the flexibility of the tree, different types of motion could be achieved. A certain percentage of the leaves could even be told to blow away when the wind reaches a particular strength. Techniques like these were actually used by MPC to create the animated 'Whomping Willow' tree in *Harry Potter and the Prisoner of Azkaban* (2004).

There are many variations of procedural animation, some of which are described overleaf.

ABOVE: The Moving Picture Company used procedural modelling and animation techniques to create small branches and several million leaves for the Whomping Willow in *Harry Potter and the Prisoner of Azkaban* (2004).

LEFT: To create digital cloth for Yoda's clothes, ILM used procedural tools that allowed materials to flow naturally according to the character's movements. *Left:* Animators could view the cloth's automatically generated movement and adjust it manually to refine its performance. *Right:* Yoda in his computer-generated clothes in *Revenge of the Sith* (2005). The procedurally generated cloth texture included subtle fluff and burrs.

PARTICLE SYSTEMS

Procedural systems are particularly good at simulating the look and movement of moving materials such as dust, snow, smoke, water or even plagues of locusts. In such cases 'particle systems' are used to create and control large numbers of objects that would be impossible to handle using key-frame animation.

To use a particle system, a computer animator defines one or a number of 'emitters' – that is, points from which particles will emerge on screen. These emitters can be imagined as invisible fireworks from which thousands of sparks emerge when lit, although they can actually be of any size or shape and can themselves be animated to move.

The animator can define how fast particles are emitted, how many to have, how they are emitted, how far and fast they travel, how quickly they disperse, how they are affected by wind and gravity, whether they grow in size, and so on. After these parameters have been set the computer will produce a simulation, normally using a low number of particles, to show the animator what the flow looks like. The parameters can then be adjusted until the desired effect is achieved. During final rendering the particles can be rendered with shaders that make them look like dust, water, fire or any other material, or they can be replaced with objects such as pre-animated fish or insects, for example.

For the action-adventure movie *XXX* (2002) director Rob Cohen wanted a spectacular avalanche with a distinct personality to chase a skier down a mountainside. 'We needed to create a very controllable avalanche that we could direct according to Rob's very specific requirements,' recalls Digital Domain visual effects supervisor Joel Hynek.

After live-action plates were shot on locations around the world, Digital Domain's artists created accurate digital replicas of the mountainside topology down which the avalanche needed to flow. 'The first step in producing the avalanche was to design the way in which it would move – its overall speed and shape,' says Hynek. 'We did this by key-frame animating a spline (<158) that moved down the mountain, following the skier and conforming to any geographical features that might affect its progress. This moving line represented the leading edge of the avalanche. When the director was happy with that motion we had to create the snow itself.'

The bulk of the snow was produced by the animated spline, which was used as a particle emitter. The spline spewed out hundreds of thousands of particles that shot out in front of it as it travelled down the mountain. The particles moved at a slower rate than the spline itself so that a massive cloud of flowing particles built up behind the leading edge as it travelled forwards. Particle movement was programmed to be affected by the downward gravity of the slope and the pressure of the wind that built up in front of the avalanche and flowed over the top.

Rob Cohen wanted many tendrils of snow to leap out of the avalanche like arms grabbing for the escaping skier. 'We selected areas of the mountainside from which we wanted these fingerlike tendrils to emerge,' recalls Hynek. 'We then created the equivalent of a minefield, with lots of particle emitters lying dormant under the snow. When the leading edge of the avalanche went over these mines they exploded high-velocity plumes of particles that shot up out of the main body of snow. Some of these emitters birthed 3-D chunks of ice. The chunks themselves were also particle emitters so that a trail of snow streamed behind them as they flew through the air. When the projectile chunks of ice hit the ground in front of the avalanche another particle emitter would automatically generate a smaller impact plume.'

The clouds of snow particles were finally rendered using volumetric rendering techniques (236>) to produce dramatic clouds of billowing snow and ice that were composited into the live-action footage.

RIGID BODY DYNAMICS

Rigid body dynamics is an advanced form of particle animation that is used to simulate the motion of three-dimensional 'solid' objects as they collide and interact. This technique is usually used to create shots that show the destruction of buildings and machinery.

Using established laws of physics and mechanics in the form of highly complex algorithms, rigid body dynamic simulations calculate the way that objects of irregular shape, size and mass will move as they travel through 3-D space. For example, a triangular-shaped object that is heavier at one end than the other will fall with a unique tumbling, rotational action that would be hard to animate accurately by hand. Because computer-generated objects are not 'solid' and can pass through one another, the process also uses collision detection methods to recognize the time and place at which moving objects touch one another. It then calculates the forces at work as those bodies meet and sends them in other directions with the appropriate motion and velocity, as if they had actually been solid objects bouncing off one another. This technique is highly processor-intensive and so is normally only used to animate the most complex of shots.

Rigid body dynamics have been used for spectacular destruction scenes such as the explosive buckling of the USS *Arizona* in *Pearl Harbor* (2001) and the apocalyptic, twisting destruction of a road bridge in *War of the Worlds* (2005). The technique was also used to generate shots of the vast tower of Barad-dûr collapsing at the climax of *The Return of the King* (2003).

Most shots of Barad-dûr were achieved using a beautifully sculpted 8 m (26 ft) high model created by Weta Workshop. But trying to create a dramatic, controllable and repeatable performance from a collapsing miniature would have proved difficult. It was therefore decided to achieve the shot digitally. The first step was to create an extremely high-polygon digital model of the tower. If a digital model needs to divide into many pieces during its destruction then it must be constructed from a large number of polygons, allowing plenty of places for breaks to occur. The finished model of Barad-dûr, created by Weta Digital, comprised an incredible 250 million polygons.

In order for the tower to fall apart the points at which it would crack had to be defined. 'We used a 2-D particle system on the surface of the tower,' explains Jim Hourihan of Tweak Films, who wrote the simulation software for this sequence. 'We basically set a bunch of particles loose on the surface and they bounced around in all directions, drawing a line behind them wherever they went. The result was an elaborate and completely random jigsaw of lines where the surface of the tower would fracture. We then used a similar method to send particles into the actual body of the tower, dividing it up into 3-D chunks.'

Next the actual collapse of the tower had to be achieved. 'We did try doing it entirely according to the laws of physics,' states Hourihan. 'We essentially switched off the glue that was holding all the bits of tower together and let it fall according to our rigid body algorithms. But the result, however realistic, just didn't look spectacular enough. The thing about animating procedurally is that you need to build in lots of control. Naturalistic simulation is often not dramatic enough for the movies and directors always want to say exactly how things should look – whatever the math says. You therefore need to maintain control of as many variables as possible in order to influence the outcome of any simulation.'

To create a more spectacular collapse sequence, the tower had a large 'box' placed around it. This box worked like an invisible force field that held the tower together. By animating the box downwards so that the tower was gradually revealed, the building could be made to collapse from the top down. 'As the tower started to crumble, additional boxes were placed around the falling clusters of chunks,' explains Hourihan. 'As these larger chunks fell the boxes holding them together could also be animated to move around, allowing smaller sub-chunks out of their influence and letting them drop away independently.'

Using this technique, the crumbling of Barad-dûr was choreographed to look as spectacular as possible, with tens of thousands of chunks tumbling to the ground, realistically bouncing off one another on their way down.



To create an awesome avalanche for *Stealth* (2005), Digital Domain added computer-generated snow to a live-action plate containing a real skier.

a: The skier was first filmed on location.

b: After a CG model of the mountain was built, an avalanche of small snow particles was created using particle animation and volume rendering.


c: Larger chunks of ice were also animated using similar techniques.

d: The elements were composited together to create an exhilarating and awesome scene.



a

b



c



d

THE LORD OF THE RINGS

After J.R.R. Tolkien's books were published in the 50s, filming *The Lord of the Rings* proved an unachievable goal for a number of notable film-makers. In 1978 the first half of the story was finally realized as a cartoon made by Ralph Bakshi – animation the only technique seemingly capable of portraying the saga's mythic characters and epic events. But it wouldn't be until the age of computer-generated imagery that Tolkien's dreams would finally be rendered with a brilliance to match their creator's imagination.

After several false starts Peter Jackson (<44) finally began shooting *The Fellowship of the Ring*, *The Two Towers* and *The Return of the King*, back to back, in 1998. Based in Wellington, New Zealand, Jackson's own facilities Weta Workshop (handling models, props and make-up) and Weta Digital (creating digital effects) worked for six years on the three films. In that time an astonishing array of techniques both old and new were adopted, adapted and combined to create some of the most awesome visual effects imagery ever seen.

Stunning New Zealand locations and full-scale sets were augmented by dozens of extraordinary miniatures, some so big that they were dubbed 'bigatures'. The regal city of Minas Tirith, for example, was 7.5 m (25 ft) high and 12.25 m (40 ft) wide with over 1,000 exquisitely detailed buildings. Once filmed the model was populated by digital characters and placed in a 3-D matte-painted mountain environment.

The trilogy's vast battle scenes were created using a custom-built artificial intelligence animation system called 'Massive'. In some shots up to 220,000 humans, orcs and elves fought according to a complex set of rules that told each how to move, react, fight and die.

A number of superb CG characters were created with absolute realism. These included the magnificent flame-spewing Balrog, the squidlike Watcher in the Water, herds of elephantine Mûmaks and the trilogy's signature achievement: Gollum. The new technique of subsurface scattering was used to render Gollum's photorealistic skin while actor Andy Serkis provided voice recordings and a motion-captured performance that was the basis of perhaps the first truly emotionally expressive CG character ever created for a film.

The films' central characters needed to vary in height from 1.25 m (4 ft) hobbits to 2 m (6 ft) wizards. To achieve this using normal-sized actors, a number of techniques were employed. For long shots, tall or small stand-ins wearing rubber masks made the main characters appear to be the right scale. Where large and small characters needed to interact, in-camera forced-perspective techniques were used – large characters like Gandalf being near the camera and small characters such as Frodo being further away. Deep-focus photography made the separated characters appear as if next to one another. Performers were also motion-control filmed in front of blue screens and adjusted for size during compositing. Many sets, props, plants and animals were also required in two sizes to make actors appear comparatively large or small.

Overseen by co-founder Richard Taylor, Weta Workshop produced thousands of elaborately detailed miniatures, costumes, weapons and props. The film's make-up requirements were also handled by the workshop, creating everything from full-body Uruk-hai warrior costumes to thousands of silicone hobbit feet.

Finally, the films were digitally graded, allowing the colour and tonal quality of every scene to reflect its place in the drama, from starkly threatening battle scenes to richly autumnal elven valleys.

After working tirelessly for eight years, Peter Jackson had created one of the most remarkable movie series ever produced. In terms of the range of techniques and quality of work, the movies were perhaps the most satisfying visual effects films for decades. The effort was rewarded handsomely at both the box office and the award ceremonies, with worldwide takings of \$3 billion and a final count of 17 Academy Awards.



WATER

Water has traditionally been a problematic substance for special effects artists when trying to film miniature ships and floods. Perhaps not surprisingly it has continued to be one of the most demanding substances to replicate digitally.

The first successful digital water shots were of wide, unbroken surfaces such as the gently rolling seascapes in *Waterworld* (1995) and *Deep Blue Sea* (1999). The water for these films used a system pioneered by Areté Image Software that used oceanographic research to produce a model of how water surfaces move under varied conditions. By altering influences such as the speed and direction of the wind, an accurate animated model of the ocean surface could be produced procedurally. However, these early seas could not have breaking waves or interact with objects to produce foam or spray. The wake seen behind the ship in *Titanic* (1997), for example, was created by filming the water around a real ship and mapping that footage on top of digital water.

The movie widely recognized as the breakthrough for realistic digital water was *The Perfect Storm* (2000), for which ILM created dramatic shots of storm-tossed seas with naturalistic foam, spray and ships' wakes.

'Making CG water look real is always going to be tricky,' states Jim Hourihan, head of research & development at Tweak Films, where a number of spectacular flood shots were created for the environmental disaster movie *The Day After Tomorrow* (2004). 'Water just has so many natural facets that need to be spot on. And even though none of us has ever seen a giant tidal wave like the one in the movie, it's one of those things that people will feel looks fake unless everything is just right.'

Tweak Films was asked to create several shots showing vast quantities of water crashing through Manhattan. The first stage of the process was to produce an accurate simulation of how a massive body of water would flow through the city's streets. 'We used a physics-based fluid dynamics simulation in which the natural properties of water are defined as an algorithm,' explains Hourihan. 'This algorithm is a mathematical description of how water moves and reacts to various forces. As a result, we essentially have to tell the computer how large a volume of water we are dealing with and what environment that water is in. We then ask the computer to calculate exactly how that water would flow through the environment.'


One of the key shots in the flood sequence was an aerial view of the tidal wave swamping the New York public library. Tweak used an accurate digital model of the library and its environs that had been created for the production by Digital Domain. The model would eventually have highly detailed texture maps applied to its surfaces to create a photorealistic view of the city, but while working on the water the model contained only the basic geometry with which the water would interact. Once the computer had been given information about the real-life scale of the city environment and the volume of water required to flow through it, a simulation was run to calculate how the wave would look and behave.

The result of the initial simulation was not particularly exciting viewing. 'What we ended up with looked like a big blob of thick blue paint flowing past a bunch of grey boxes,' says Hourihan. 'It was quite hard to see how it would eventually look like a massive tidal wave because it didn't have any splashes or surface details that might indicate scale. It was just a blob. Nevertheless, it was a blob that was moving according to the laws of physics.'

Though accurate, the initial simulations didn't perform in the way that the film's director, Roland Emmerich, wanted. 'We had to do lots of simulations before the water behaved in the desired dramatic fashion,' says Hourihan. 'We influenced its performance by pumping more water into the simulation as it flowed along. We could make it swell or leap up by injecting more water at the places it was needed. We were effectively art-directing nature.'

When the director was satisfied with the character of the water, one problem still remained. 'We were pumping in so much water that it was totally engulfing the library. It was just flowing right up and over the roof,' says Hourihan. After trying more simulations Tweak eventually solved the problem by making the library building artificially high. Because the camera was looking down from above, it wasn't noticeable that the building had grown dramatically for the purposes of the simulation.

With the basic character and surface topology of the water simulated, the next task was to make it look less like paint and more like water. This was largely the job of shaders (<167) which would help to create the look of a large body of water during rendering (236>). However, a number of additional watery details had to be added. 'One of the things that makes rough water look real is the foam that floats on its surface,' explains Hourihan. 'We produced a number of two-dimensional animated texture maps to simulate what that foam looked like. They were basically procedurally produced black-and-white swirling patterns. We then scattered those maps all over the surface of the water.' During the initial water simulation the computer had also generated velocity vectors describing how the surface of the water moved in relation to the main body. 'We attached the maps to the main body of water so that they would be driven by its gross movement. We then used the velocity vectors to push the foam around on the surface so that it looked like it was floating over, rather than being fixed to, the main body. It was a bit like attaching the foam to the body of water with bits of elastic – the foam went



ABOVE: To create incredible scenes of a flooded New York for *The Day After Tomorrow* (2004), Tweak Films used a high-resolution model of the city. Here (a) the New York Public Library model can be seen from above. Using procedural animation techniques, the flow of the flood waters is simulated as it engulfs buildings (b). Additional simulations create foam and spray on the top and leading edge of the water (c). The final composite is a chilling and dramatic demonstration of a city overpowered by the forces of nature (d).

TOP RIGHT: In this scene from *The Day After Tomorrow*, the steps and deep water through which people are struggling are an interior studio set. The background buildings were added later, along with a computer-generated tidal wave that smashes everything before it.

RIGHT: *The Perfect Storm* (2000) featured some of the first truly natural-looking CG water. Here, the gallant fishing boat *Andrea Gail* is dwarfed by ILM's spectacular digital tempest.



where the water did but still had a fair amount of freedom to slip and slide.' Because the animated surface textures were attached to the body of water they would eventually be pulled down into that water as it churned. 'After a few frames the textures would get submerged or buckle and turn into a grey mush. So we had to continually recycle the maps, bringing fresh new ones onto the surface as the old ones disappeared.'

As the massive wave forced its way through the city it needed to produce spray wherever it struck objects and buildings. The computer was able to detect where the leading edge of the water hit or interacted with objects and at those points particle emitters sent spray into the air (<220). The amount and force of the spray was determined by the water velocity vectors calculated as part of the simulation. The movement of the spray was also influenced by surface wind



currents that were also calculated as part of the main water simulation.

Tweak's artists also manually placed emitters wherever they felt more white water was needed. 'The director particularly wanted a very dramatic surge of spray when the wave comes around the library from two directions and meets in the middle,' remembers Hourihan. 'When we ran the simulation the two bodies of water met and there was just a bit of a ripple, but Roland Emmerich wanted a big dramatic moment. In the end we had to pump a lot of extra water under that area and then add a slightly delayed fountain of spray. It was completely unnatural but the timing and size of the splash certainly made for a dynamic moment.'

Considerable effort was put into making the spray look natural. 'If the particles were too small they looked granular and more like sand, and if they were too big they looked blobby. We eventually programmed the particles to expand until they reached the next-nearest particle – so some bits of spray were large and some were small,' says Hourihan. But the key to getting the particles looking perfect was the way they created shadow: each tiny particle of spray needed to cast a shadow behind itself to make it stand out. Without this shadow the many particles in a plume of spray would look like one homogenous blob. As the airborne particles fell back down they intersected with the main body of water, remaining visible under the surface for a while to help produce a foamy, churning look.

The main body of water was rendered using subsurface scattering – a technique commonly used for skin (<204). 'The light would go down into the water and interact with things like the submerged foam texture maps and the sinking spray particles before coming back out of the water in a number of directions with a really naturalistic quality,' says Hourihan.

The various water elements were composited with the photorealistic New York model that even included dozens of tiny vehicles being swept along in front of the wave. The result was some of the most awesomely apocalyptic images of natural phenomena yet created for the movies.




FUR AND FEATHERS

Along with water, the creation of lifelike hair and fur became a technical and creative holy grail from the earliest days of computer animation. The first sophisticated digital characters could only have smooth or solid outer surfaces, such as the metallic T-1000 character in *Terminator 2: Judgment Day* (1991), or the dinosaurs of *Jurassic Park* (1993). Creatures with a furry exterior, such as the computer-generated penguins in *Batman Returns* (1992), only looked effective in long shots – their appearance being the result of shaders (<167) that made light hitting them look as if it was bouncing off a furry surface.

With improvements in processor power and advances in modelling, animation and rendering, techniques progressed from producing the illusion of a hairy surface to the actual modelling of individual hairs themselves. An early landmark in computer-generated fur was *Jumanji* (1995), for which ILM created a digital menagerie with varying degrees of success. More recently, films such as *Harry Potter and the Prisoner of Azkaban* (2004) and *The Lion, the Witch and the Wardrobe* (2005) have proved that, though still a daunting challenge, the creation of convincing photorealistic hair, fur and even feathers is no longer an obstacle.

'Computer-generated fur has traditionally been one of the hardest things to create for two main reasons,' explains Ben Morris, CG supervisor at London's Framestore CFC, where stunningly realistic digital squirrels were created for *Charlie and the Chocolate Factory* (2005). 'Fur has many natural qualities that make it complicated to replicate digitally,' expounds Morris. 'The way it responds to light – reflecting, absorbing and shadowing itself, its colour, condition and distribution and the way it reacts dynamically to the movement of the creature – these are all aesthetic challenges that must be addressed to produce a

ABOVE: Along with water, simulating realistic fur has been one of the greatest challenges for digital artists. This early CG monkey was created by ILM for *Jumanji* (1995).



realistic result. The other complicating factor is the sheer amount of data involved. A single creature may have many millions of hairs, which traditionally created extremely large files and long render times.'

For *Charlie and the Chocolate Factory*, two dozen real squirrels were painstakingly trained for a scene in which they sort nuts before attacking one of Willy Wonka's young guests. Many tasks were too complex to achieve with live squirrels, however, so computer-generated rodents had to be used for over 60 shots. 'Charlie was an immense challenge because our CG squirrels had to perform in parallel with real creatures – often very close to the camera and frequently cutting back and forth between live and CG squirrels,' states Morris. 'This meant our squirrels had to be absolutely indistinguishable from the genuine thing. From the start we knew that the success of our squirrels was going to depend on getting their fur to look right. We also had to find a way to do it efficiently. With almost 100 CG squirrels in some shots the amount of data would quickly become unmanageable without a highly optimized production pipeline.'

Morris and a team of 20 modellers, animators, technical directors and composers spent eight months studying real squirrels and devising the systems and tools needed to create their CG doubles. Digital squirrel models were built using standard modelling techniques to create their skeleton, musculature and outer skin. 'Our squirrels were physically identical to the real thing except for some slight modifications to their paws,' says Morris. 'In nature a squirrel's hands are designed for holding nuts and climbing trees – but ours also had to be capable of roughing up small girls!'

With the basic body complete, attention turned to the all-important coat of fur. 'When you stare at squirrels all day you quickly realize that they look as different from one another as humans do,' says Morris. 'Many of these differences are due to their fur, which can have various patterns and colours, be fluffy or oily, flat or spiky, clumpy or very fine. These differences can be due to a squirrel's age, sex,

ABOVE: To create the distinct appearance of genuine squirrel fur for *Charlie and the Chocolate Factory* (2005), texture and colour maps like these were applied.

Top row:
 (left) alterations to the angle of hairs
 (centre) basic colour
 (right) overall fur length

Bottom row:
 (left) distribution of guard hairs
 (centre) distribution of hair curl
 (right) tail-fur length



condition or even mood. At first we thought we would just be making one generic squirrel but in the end we created 26 variations – each of which we named after musicians that they resembled. David [Lee Roth] was a squirrel with a bit of an 80s hairdo, Iggy [Pop] was really scrawny with a bit of mange and Herbie [Hancock] was darker with heavier, clumpy fur.

Morris and his team established that their squirrels would need seven different types of fur for the various areas of their body. Each squirrel had back fur, face fur, hand fur, belly fur, tail fur, underfur, and guard furs as well as whiskers and eyelashes. Each of these varied in length, thickness and colour and had to be distributed and intermingled in the correct way to look convincing.

To create its covering of fur the outer skin of the squirrel model was unwrapped to produce a flat template showing where the creature's legs, eyes, ears and other features were located. 'Painting on copies of the unwrapped skin we created a series of maps (<164) that would describe the various quantities and qualities of fur on different parts of the squirrels' bodies,' explains Morris. 'The maps were usually created in greyscale – the lighter the map the thicker or longer the fur in that area, for example. Separate maps defined the direction that the fur pointed, or places where it might be more clumpy. Other maps identified the seams around the animal's joints where extreme movements caused the hair to form a parting right down to the skin below. There was also a basic colour map that described what the squirrel's skin looked like at the base of the fur. In the end each squirrel had around 70 separate maps to control the look of the fur all over its body.' Additional maps also ran the length of each strand of hair. 'Squirrel hairs change colour several times between their base and tip,' says Morris. 'The most obvious example is the tail hairs. If you look at a squirrel's tail end-on you will see many concentric circles of colour changing from whites and creams to blacks, greys and browns.'

The fur itself was made up of thousands of simple geometric primitives – each piece being several vertices connected to form a curve. These hairs had no breadth or other physical characteristics. 'Our squirrel hairs didn't have a circular cross-section like real hairs do,' says Morris, 'They were actually completely flat. Whatever angle the camera viewed the hairs from, they always appear round because the shaders used during the rendering process instructed light to bounce off them in a way that made them appear cylindrical.'

Each squirrel was populated with approximately five million hairs, a quantity chosen through visual experimentation rather than by counting the actual number of hairs on a real squirrel. As a model covered in so much fur would be impractical to work with, most hairs were only ever generated



LEFT: Computer-generated squirrels used for close-up shots in *Charlie and the Chocolate Factory* (2005) were covered with as many as 5 million hairs. As the squirrels moved away from the camera, their coats were automatically 'pruned' until they comprised as few as 250,000 hairs.

BELOW: A contingent of computer-generated squirrels checks the quality of Willy Wonka's nuts. For some shots such as this, CG squirrels were used in the foreground, while background squirrels were animatronic puppets.



when the final squirrel images were rendered. Before then each squirrel was covered with approximately 2,000 'guide' hairs. 'Most fur movement was driven by the way the underlying skin to which it was attached travelled over the body as the squirrel moved around,' explains Morris. 'Even so, we had to continually groom the hair to make sure it looked at its best. Each of the two thousand guide hairs was programmed to influence the thousands of as yet invisible hairs around it. After each shot of a squirrel had been animated we would render a few final frames to check how the full fur coat looked and then rearrange any unsatisfactory areas by physically tweaking the nearest guide hairs. This hair-dressing took an immense amount of time and patience.'

Essential to the look of the final squirrel fur was the way it was lit. 'Any photographer will tell you that hair and fur is tricky to light,' says Morris. 'A common trick for making fur look good is to add a bit of backlighting to create a rim-lit feel which nicely shows off any detail and texture in the individual hairs.' Unfortunately the scene in which the squirrels were to appear took place in a large room with uniformly diffuse lighting. 'It was about the hardest environmental conditions exactly, but in this case we often had to cheat the lighting of our models slightly to make them look their best.'

Another method used to create subtle lighting effects on the squirrels' fur was a pre-rendering process called 'ambient occlusion.' 'Getting accurate ray-traced [238>] shadows on the squirrels' fur during the final rendering process would have been an extraordinarily time-consuming process given the vast number of hairs contained in each image,' states Morris. 'Instead we rendered a pass for each shot in which the squirrels had no fur at all. During this render every point on the body of the animal sent out many rays in every direction in order to detect whether there were any other objects in the vicinity. Depending on how many of these rays hit another surface and how near that surface was, that area of the squirrel's body was given the appropriate amount of shadow darkening. The result was a scene in which the parts of the squirrel that were "occluded" – such as the self-shadowing areas inside the ears, mouth or armpits, for example – were grey or black while everywhere else was white. The ambient occlusion pass contained no detail other than these occlusion shadows. During the final render and compositing process this shadow information was applied to the squirrel's fur to produce very naturalistic, soft and subtle shadowing. It's a technique we use for all our animated characters and environments.'

With the appearance of the squirrel fur perfected it only remained to ensure that the production pipeline was as efficient as possible. 'The fact is that with enough time, technology and creative talent, any visual effect is now possible using computers,' remarks Morris. 'The real trick is making the process efficient.' One of the ways Framestore CFC made their squirrels faster to handle was to regulate the amount of fur on each squirrel according to its proximity to the camera. 'There's no point having five million time-consuming hairs on a squirrel when it's way off in the distance,' says Morris. 'For this reason we designed a system that regulated the number of hairs [level of detail, or LOD] on a squirrel depending on how near or far it was from the camera. A squirrel up close would have all five million hairs but as it ran away hairs would automatically disappear from its body. Meanwhile the remaining hairs would grow much thicker to fill in the gaps. At their furthest distance squirrels would have fewer than 100,000 hairs, but they were so thick that seen in detail they looked more like porcupines!' Another time-saving trick involved identifying which hairs were out of sight of the camera and then instructing the computer not to generate that fur data while rendering the image. 'If you could sneak a look at the other side of our squirrels while they are performing on screen you would see that the areas pointing away from the camera are completely bald!' reveals Morris.

Squirrel shots were finally rendered in up to 20 passes to provide compositors with the ability to alter their look almost infinitely as they placed them into scenes with actors and real squirrels.

PRE-VISUALIZATION

The complexity and expense of modern movie production means that directors now spend increasing amounts of time carefully planning their shots in advance of any filming.

Traditionally, storyboards have been a popular way of planning how a movie will look. Working closely with the director, artists produce drawings that illustrate the composition, action and camera movement of every shot in a sequence. Storyboard artwork can even be shot on video and then edited together with music and dialogue to create a more accurate sense of a scene's intended pace and movement. However, these moving storyboards, called 'animatics', are now being superseded by detailed sequences of 3-D animation that are created entirely within the computer. This process is called 'pre-visualization', or 'pre-viz'.

'Pre-viz is the ultimate planning tool for directors who are making complicated movies with a lot of action or effects,' states Colin Green, president of Pixel Liberation Front (PLF), which has created pre-viz sequences for movies such as *Minority Report* (2002), *The Matrix Reloaded* (2003), and *I, Robot* (2004). 'We start by creating correctly scaled digital models of everything that will be in a scene, such as props, scenery and vehicles,' explains Green. 'We also produce animated characters that resemble the actors in the movie. For *Superman Returns* [2006] we even had a Superman character with a fluttering cape. Once we have all our elements we then work with the director to produce animated sequences based on storyboards that have already been approved.'

Traditional storyboards are designed to capture the movement and action of a shot in one or two dynamic drawings but they are rarely truly representative of the way a scene will look once filmed. 'With pre-viz directors can see exactly what they are hoping to get in the final version of the movie,' states Green. 'For a start, the actors, props and camera will be moving, which makes a tremendous difference. Then there is the ability to constantly try new ideas; if the director doesn't like the way things look we can quickly try something else. We could move the actors in a different way, try another camera angle or lens, or perhaps a different camera move. Or we can see if the

scene will work better if it is edited differently. Most pre-viz is done before shooting begins but we can also be on set during filming to quickly test new ideas before the real cameras and lighting equipment are set up for a new shot.'

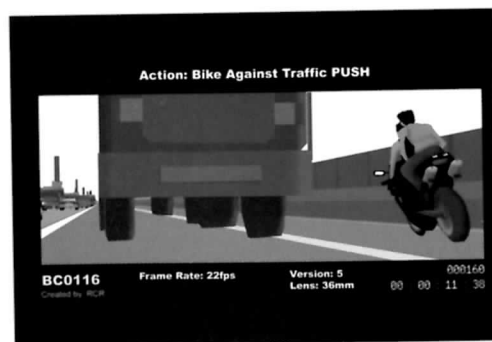
Pre-viz tends to be used for the more action-oriented scenes in a movie and is used to solve technical problems, help plan complex stunts, and aid the design and conceptualization of visual effects.

'With a sequence like the freeway chase in *Matrix Reloaded* we were able to streamline the whole of the production process from a technical standpoint,' claims Green. 'For example, we could work out the best angle to shoot the action from and what lens to use for maximum impact. We could determine in advance how fast action vehicles needed to travel and how the stunts should be performed. We could also help work out what could be done as real stunts and what would have to be done as visual effects.'


One of the greatest advantages is in the planning of visual effects sequences. 'It's often hard to imagine what visual effects will look like when completed, but with pre-viz everyone involved can see how a shot is intended to work,' says Green. 'When live-action performances have been shot against green screen we can quickly composite them into our pre-viz shots to see if they work, right there on the set. We can also extract data from our pre-viz sequences for use during the actual production of visual effects. The most useful of this information is perhaps camera movement data. Once we have produced a virtual camera move for a pre-viz shot, that data can be exported and used to program motion-control cameras being used to film live action or models. It can also be applied to the virtual camera used to film any additional CG elements.'

Pre-viz has a trickle-down effect on all departments involved in a film. Because a pre-visualized sequence is an accurate blueprint of everything in a scene, art directors can determine exactly how large they need to build their scenery, cinematographers can plan how much lighting equipment may be needed, and producers can even work out how many extras they will need to hire to populate a scene.


A completed pre-viz sequence is normally given to a movie's editor so they can cut it into their working edit of a film. As filming progresses, pre-viz shots will gradually be replaced with live-action footage and completed visual effects shots until the finished sequence is achieved.



a



b



c

To plan complex action scenes, animated pre-viz graphics similar to these are created by Pixel Logic.

- a: Original pre-viz animation using basic models to plot action and camera movement.
- b: Layout of pre-viz animation as viewed from above and used to plan the movement of real vehicles and cameras during filming.
- c: Exact plotting of all elements in a shot on a frame-by-frame basis, including additional detail such as compass directions to help plot direction of the sun.

DIGITAL LIGHTING

Just like shooting a movie in the real world, digital scenes must be properly lit before filming takes place. 'A model built in the computer exists within a totally dark environment', explains Craig Ring, visual effects supervisor at DreamWorks Animation, who has overseen lighting and effects in animated films including *Antz* (1998) and *Over the Hedge* (2006). 'During the animation process, a very flat type of light is used to view scenes in the computer. Without adding further specialized lights to the scene, it would be in total darkness when finally rendered. Every form of light and lighting effect you see in a CG scene is therefore designed and created by us.'

Much like cinematographers on real film sets, digital lighting designers have a full range of lighting tools at their disposal. The designer can choose from spotlights, which provide focused cones of light that cast shadows; radial lights, which cast light outwards in all directions like a real light bulb; ambient lights; non-directional diffuse lighting that illuminates a scene uniformly; and global lights, which cover the whole scene with parallel rays of light from a distant spot, much like the sun.

'As a visual aid the lights we use in our scenes are designed to look like real lights,' explains Ring. 'We have little icons of fluorescent tube lights, and spotlights that even have movable barn doors on the front. It's just like having a miniature sound stage on the screen in front of you. These lights don't behave exactly like the real thing, but they are our nearest equivalent. When we want a light, we can pick up one of these icons and drag it to the place in the digital environment where we want to use it – just as a stage technician would pick up a light and place it where they want before plugging it in.' Digital lights themselves are invisible; only the light they produce can be seen, so they never have to be concealed behind props or scenery. Once a light has been positioned within the virtual environment, the lighting designer adjusts variables such as focus, fall-off, colour and so on.

The computer environment is very different from the real world in which light is automatically affected by how shiny a surface is or by how much dust is in the atmosphere. In particular, digital lights do not automatically produce any diffuse interreflectivity – the effect of light altering in quality as it bounces between objects – as Ring explains: 'In a real environment light hitting an object will automatically change colour and become more diffused when it is reflected off again. That reflected light will then hit other objects, where its colours and characteristics are changed even more. Digital lights don't automatically work like that, so we have to find ways to mimic nature.'

Until recently the only way to achieve such effects in a digital environment was to manually place lights that artificially reproduced a similar look. In the real world, for example, a white light pointing at a red brick wall results in a small amount of reflected soft red light illuminating other objects nearby. To fake this effect lighting designers will first place their white light pointing at the wall before placing a small red light within the wall to cast some 'reflected' light away from it. This can mean placing dozens of tiny lights in every scene in order to simulate the effect of natural light interactions.

It is now possible for physically accurate lighting interactions to be automatically calculated by the computer during rendering using a technique called 'global illumination' (237>). 'Global illumination mathematically calculates how light bounces around in a scene. We just have to place the main key lights and the bounced light is calculated and added by the computer. At DreamWorks we first used these techniques extensively for *Shrek 2* [2004] and it can look very naturalistic,' says Ring. 'The problem is that the calculations can take a very long time during rendering, which is expensive.'

To save rendering time for shots using global illumination smaller objects in a scene can be 'switched off' so that they don't become part of the calculations. A shot of a room will have all the small objects used as set dressing switched off, so that only the large important objects like the furniture, walls and floor are included in the lighting calculations. 'The other thing we do is create proxy objects,' comments Ring. 'If a character is standing on some grass by a tree we would want diffuse light to be reflected onto the

character from the green leaves of the tree above and the grass below. Calculating the way light bounces off thousands of leaves and blades of grass would be impractical so we can build a basic replacement tree that is just a green blob on a brown stick, and grass that is just a flat green surface. During rendering the light reflected off these and onto the character will be quick to calculate and we will have a character with natural-looking lighting on their skin. We then composite the shot of the character back into a shot of the trees and grass that has been rendered in the normal way.'

The ability of the computer to calculate realistic lighting effects means that digital lighting designers can spend less time trying to mimic nature and more time on the finer details of a shot. 'We can be much more like real-life cinematographers now that we don't have to spend so much time compensating for the inadequacies of the digital environment,' says Ring. 'We can put more effort into using light to tell the story, to direct the viewer's eye, to make sure characters read well against the background.' Like cinematographers in the real world, digital lighting designers need to make the stars of the film look as good as possible. 'When lighting characters, there are certain small things that are important to get right,' explains Ring. 'Real cinematographers are always careful to get a bright highlight in the actors' eyes – a little sparkle that helps to bring them to life. We do exactly the same thing, though luckily for us we can animate our lights to move along in front of a character and to light only the eyes, so the sparkle is always there.'

Being able to tell a light exactly which objects to illuminate is one of the great advantages of digital lighting. In the real world, cinematographers may spend hours arranging lights so that particular actors or areas of a set have just the right quality of illumination. In the digital world, a light can be told to light just one object and nothing else. Lights can even be given a negative value to take light away from objects that are overbright. 'The flexibility of digital lights is something that film and stage lighting people get really jealous of,' exclaims Ring. 'Although we sometimes long to be able to place lights and automatically get all the lighting interactions that you do in the real world, the fact that we can animate the position of a light so that it follows a character around, or ask a light to illuminate the tip of a nose and nothing else, or drastically change the quality of the light at any time, is really very liberating.'



ABOVE: In this scene from *Shrek 2* (2004), the main lights, displayed as yellow icons, are direct lights, such as spotlights and point lights. All bounced 'fill' lighting was created using global illumination, without which the number of lights needed for this scene would have been far greater.



HIGH DYNAMIC RANGE IMAGES (HDRI)

Lighting entirely computer-generated scenes is a complex task, but lighting digital models that are to be placed within a real-world environment adds an additional level of complexity. Only when the lighting of digital models matches that of the filmed environment can the two elements be merged seamlessly.

Re-creating real-world lighting conditions in the computer is traditionally achieved by taking careful note of the intensity, position, and colour of the lights used during filming of the live-action plate. If a shot is filmed outdoors, the height and direction of the sun is also recorded. A highly reflective chrome sphere is often photographed in the environment, the reflection revealing the relative position of lights in the scene. A matte grey sphere may also be photographed to record the colour of the light. By studying these images the digital lighting artist will attempt to manually re-create the same lighting conditions within the computer.

For incredibly realistic results, the exact lighting of an environment can now be recorded photographically and used to directly generate the lighting used for a digital model. This technique is known as 'image-based lighting' and the photographs it uses are called high dynamic range images (HDRIs).

An HDRI is a high-quality panoramic photograph of the environment into which a computer-generated object is to be placed. Normal photographs are taken at an averaged exposure level to produce a balanced image. In these images most objects will be evenly exposed but areas of deep shadow will be completely black (underexposed) and contain no image information, while very bright areas will be white (overexposed) and will contain no image information. An HDRI image, however, is produced by photographing an environment several times, each time at a different exposure level. The result is a series of identical photographs that range from very overexposed to very underexposed. These images are then digitally combined to produce a single image containing a complete range of exposure detail, with information in both the brightest and darkest areas of the picture. This is an extremely accurate record of even the subtlest levels of light and colour within an environment.

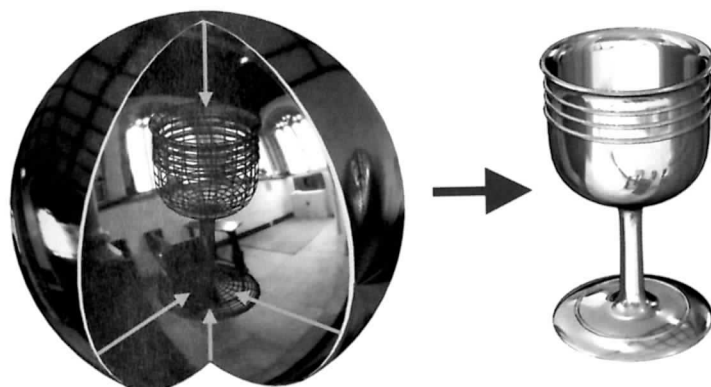
Normal photographs only show a small section of an environment as seen from a single point of view. In order to accurately re-create the lighting in a scene, HDRI images must be able to show what the entire environment looks like. To create a complete image, HDRIs can be photographed with a 180° 'fish-eye' lens or by using a normal lens to photograph reflections from a 100 per cent reflective chrome ball. Either way, the result will be a number of photographs that collectively show a complete but distorted view of the environment. Using specialized

software these images are unwrapped and stitched together to create a single HDRI of the entire live-action environment. This image is called a 'light probe'.

The light probe image is then mapped onto the inside of a sphere that surrounds the digital model that needs to be lit (fig. 9). That sphere, with all of its information about the colour, position and intensity of light in a scene, then itself becomes one large light source that illuminates anything placed within it according to the HDRI information on its surface.

HDRI lighting is a complicated process but it produces reliably accurate lighting for any computer-generated object that needs to be placed into a live-action scene. It is particularly useful for highly reflective objects since they will reflect a realistic image of the surrounding scene, even though they were never actually present during filming.

FIGURE 9 HIGH DYNAMIC RANGE IMAGES (HDRI)



Once an HDRI has been obtained, the image is 'wrapped' around a scene and used to 'project' its environmental lighting information onto objects within that scene.

THE VIRTUAL CAMERA

FAR LEFT: To create a High Dynamic Range Image (HDRI), a mirrored ball is photographed in the desired location. Several photographs will be taken at different exposures in order to capture the full range of light and shadow.

LEFT: The spherical photographs from the mirrored ball are unwrapped and used to produce an image which becomes the basis for the lighting of any CG objects placed in the environment.

The 'virtual camera' is the device used to 'film' the world inside the computer. The digital domain does not, of course, actually exist as a physical location that can be photographed in any traditional sense. However, the virtual camera is used to describe the viewpoint that the computer will use during the final image-producing process of rendering (236>).


Live-action motion picture cameras are constrained by their physical size and weight as well as the manual methods used to adjust their aperture, focal length and focus. These constraints have created a series of visual storytelling conventions with which audiences are familiar and that they accept as the norm when they watch a movie.

The virtual camera suffers from none of the constraints of its real-world equivalent. The virtual camera does not actually exist as an object, so it has no size or weight – it can go anywhere, at any speed. Because it has no optical lenses and holds no film there is no need to worry about the amount of light in a scene or how that affects aperture, film speed, exposure or depth of field.

However, because audiences are so familiar with the look produced by the traditional movie camera, the virtual camera has had many of the characteristics of a real camera imposed upon it. 'The camera that we use to film our computer environments has been designed to replicate real cameras,' explains Damon O'Beirne, head of layout for *Over the Hedge* (2006) at DreamWorks Animation. 'We can give our camera any focal length lens and it will exactly replicate the look produced by that lens when used on a real camera. We can zoom, pan, tilt, dolly – all the moves that a normal camera can do. We can also do a lot of things that a normal camera can't do.'

Since the virtual camera has no constraints, it can be used to produce breathtaking rollercoaster rides through the artificial environment of the computer. 'That's fine in certain circumstances,' says O'Beirne, 'but we tend to use a more traditional style of cinematography that is more grounded in the real world. We

FIGURE 10 THE VIRTUAL CAMERA



often watch movies shot on film and select a cinematic style that we will try to emulate. For *Over the Hedge* we used *Raising Arizona* [1987] as our main cinematic reference. It's important to impose a distinct style because instead of a single cinematographer as on a normal movie, we have half a dozen people creating the camera moves for a show. Each of these artists will have their own film-making influences – be it Steven Spielberg or John Woo – and they'll each end up making their own style of movie if there isn't a clear overall vision!

To help replicate the look of real-world photography, O'Beirne likes to limit the toolset that his artists are able to use. 'In the real world cinematographers will have a limited range of lenses in their kit – a 35 mm, a 50 mm, and so on. But our virtual cameras can have any focal length lens we want. We can just dial in 33.062 mm and that's the lens we'll get. This can lead to an inconsistency of style so I'll put together a set of lens sizes that can be used in each film. If the artists are using a 24 mm lens and they want to go a little wider they can't just dial in another half-millimeter, they have to switch to the next lens down – say a 21 mm – just like a real cinematographer would.'

Virtual cinematographers also try to create camera moves that look as if they were produced by relying on any of the standard techniques used to move a camera on a real film set. These include the gentle glide of the steadicam shot, the sweeping arc of a crane shot, the uneven drift of the hand-held camera and the swift linear moves produced by a dolly track. It's not only the quality of the movement that is replicated in the digital world, as O'Beirne admits: 'We also sometimes build in small mistakes that you might expect to see in a live-action film. For example when a character walks and our camera pans to follow them, we might delay the start of the camera movement very slightly as if there were a real camera operator reacting momentarily after the movement of an actor. We might also put in a little camera shake when a camera moves, as if it's actually travelling along a dolly track!'

Layout artists don't always stick slavishly to the restrictions of real-world photography, however, and there are plenty of times that the virtual camera's abilities are used to positive advantage. 'In the real world cinematographers often struggle to get enough depth of field, especially when the light is fading,' says O'Beirne. 'We, on the other hand, can have a night-time scene with pin-sharp focus from the front to the back of a shot, if we wish. That's pretty handy!'

Virtual cameras are animated much like any other object within a computer-generated scene (fig. 10, <233). The digital artist begins by defining the size of the lens and the number of frames that the shot in question will last. The camera is then placed within the 3-D environment. When the artist is happy with the view from the camera's start position, he or she will set the first key frame. The artist then moves the camera to its second position and, if using a zoom lens, may adjust the focal length. In the real world the camera operator must concentrate on moving the camera to follow the important action within a scene. In the world of the computer, the camera can use 'target tracking' to keep the subject automatically in the centre of the frame, however extreme the movement of the camera is. The precision of target tracking tends to produce artificially perfect shots, but the method can be used as a starting point when animating a sequence.

The digital artist builds up the movement of the camera one key frame at a time, continually running the camera backwards and forwards, studying its motion and making refinements where necessary until a satisfactory shot has been achieved and the shot is ready to be rendered.

While the virtual camera and its cinematographer are capable of replicating any shot that can be captured with a traditional camera, some directors find it frustrating that they do not have a camera that they can physically handle during photography. This lack of interaction between camera and operator has driven some traditional film-makers to seek alternative ways to create their virtual camera moves. For *The Lord of the Rings: The Fellowship of the Ring* (2001) director Peter Jackson wanted to re-create the dramatic look and feel of a hand-held camera in a computer-generated scene involving a rampaging cave troll. The scene was first pre-visualized (<230) using motion-captured performances for the troll and the characters that he was attacking. In a motion-capture studio, Jackson was then able to view the animated fight scene using virtual reality goggles. A hand-held camera was created out of a wooden block on a pole. This was fitted with motion-capture markers so that as Jackson moved the camera around the set his motion was relayed to the computer and he was able to view what it was 'filming' through his goggles. In this way Jackson was able to follow the pre-planned fight, spinning around to see interesting action and leaping out of the way when the troll came in his direction. The resulting camera movements were edited and refined before being used to film the final CG environments and characters. The result was a thrilling action sequence that has the frenetic immediacy of combat-zone news coverage.

Another method of creating real-world camera moves was developed by Sony Pictures Imageworks for the shooting of its animated film *The Polar Express* (2004). When animated character performances had been edited to create final scenes they were 'filmed' using live camera moves created by cinematographer Rob Presley. Presley used a device that was built to resemble the mechanism used to control the movement of a real camera. By manually turning a set of wheels as he watched the animated action he could pan and tilt the virtual camera to shoot the film like a live-action movie.



ABOVE: In this scene from *Shrek 2*, the virtual camera is displayed as a yellow icon on the left. The digital cinematographer programs the camera to follow the pre-animated character movement. The grey lines emerging from the camera indicate the area of the scene that will be visible on screen.

FAR RIGHT: For this complex scene from *Starship Troopers 2* (2004) created by Tippett Studio, both the environment and the movement of the camera needed to be carefully tracked so that additional CG elements could be added.

- a: The live-action background plate of soldiers is filmed inside a studio.
- b: The environment is modelled in the computer, matching the live-action set exactly. The scene is then replayed one frame at a time, adjusting the virtual camera so that it matches the movement of the real camera.
- c: The CG environment with its matched camera move is used to produce animation effects, in this case hundreds of animated alien insects.
- d: The CG creatures are tested against the live-action environment. Their exploding bodies appear to interact with the match-moved concrete barricades.
- e: The animated aliens and explosions are rendered.
- f: The rendered elements are composited with the live action to produce the final shot.

MATCH MOVING

Designing camera movements for completely computer-generated scenes gives digital artists all the freedom that real-world camera operators have to move a camera and produce the best compositions. However, many digital effects are eventually composited into footage of real-world environments that have already been filmed with a moving camera: animated dinosaurs are seen crashing through real forests, digital ships land their simulated armies on actual beaches, and computer-generated vehicles race along genuine roads. In such cases the movements of the virtual camera used to film the computer-generated elements must precisely match those of the camera used to film the live-action environment. The resulting images, created at different times and using differing technologies, will have identical camera movements, allowing them to be layered together seamlessly during compositing.

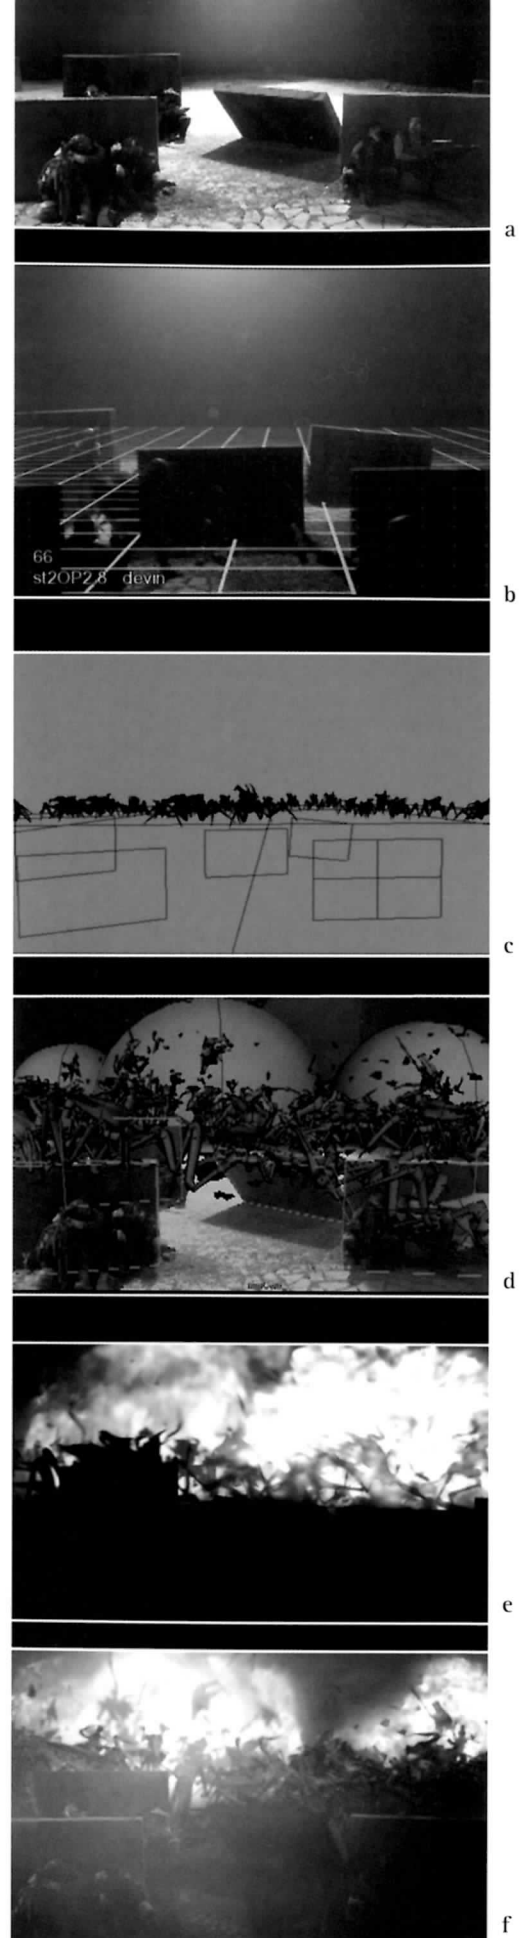
To replicate the movement of the real-world camera in the computer, a process known as '3-D camera match moving' is used. During photography of the live-action plate, special effects supervisors take accurate measurements of the position of each significant object in a scene, including specially placed markers. The starting position of the camera is also measured and details of its lens type are noted. These measurements are sometimes recorded using the highly accurate laser measuring equipment that is used by cartographers and architects. The information collected on location is used to construct a basic yet precise 3-D model of the real location within the computer. This digital reconstruction of the real environment is essential because any CG objects that will eventually need to appear to interact with the real world, for instance digital vehicles, must be animated to move over an accurate representation of the real ground plane and avoid, or possibly interact with, any of the scene's real objects.

To create a camera match-move, the footage of the live-action background plate is viewed on a monitor. The digital model of the environment is then superimposed over the background plate, and the virtual camera is manoeuvred until its view of the model is perfectly aligned with the viewpoint of the camera used to film the live action. When the first frame of the live-action plate and the digital environment are perfectly aligned, the operator moves through the shot a few frames at a time, altering the position of the virtual camera so that its perspective remains identical to that of the live-action camera. By working through the whole shot in this way, the virtual camera is programmed to produce a perfect copy of the location camera's movements.

Another method of match moving does not involve creating a model of the real environment but instead relies on identifying and tagging a series of fixed points within the live-action plate. These points might be anything that is small yet detailed or of high enough contrast to be identified in each subsequent frame, such as the corner of buildings or possibly specially placed tracking markers. With a number of points identified, sophisticated match-moving software will advance through the sequence, frame by frame, detecting the changing position and relationship of each identified marker point and extrapolating accurate camera motion data. Once the live-action camera motion data has been gathered it can be applied to the virtual camera used to 'film' the necessary computer-generated elements. The data can also be exported to motion-control cameras for the shooting of additional live-action footage or visual effects elements such as miniatures. As a result of using the same camera movement to film every element in a shot, the layers will sit perfectly on top of one another in the final composite.

Another important form of tracking involves recording not simply the movement of the camera used to film a scene, but also the movement of the actual objects within a scene. With the motion of real objects tracked they can be modified digitally. An example might be to track the movement of a real pick-up truck filmed as it drives through a shot so that a computer-generated farm dog can later be placed riding in the back. To track the motion of a real object a digital model is first made with the same basic shape and scale. This 'proxy' object is then displayed on a monitor with the live-action footage displayed behind it. The digital artist then adjusts the position of the proxy model, in our example the pick-up truck, until its perspective perfectly matches that of the real truck and the two are aligned, one on top of the other. The shot is then advanced one frame at a time and the position of the model truck is changed incrementally so that it keeps pace with the real-world version. The result in our example will be a model pick-up truck that drives in exactly the same way as the real vehicle did during filming. Our computer-generated farm dog can then be animated standing on the back of the digital truck. The shot can then be rendered without the digital truck and the resulting shot of the dog gliding along on thin air can be composited together with the live-action vehicle footage, the two matching perfectly. Interesting examples of this technique have been the replacement of the actor Ralph Fiennes's nose with a digital nose in *Harry Potter and the Goblet of Fire* (2005) and adding a digital blade to a bladeless sword handle wielded by Tom Cruise when filming fight scenes for *The Last Samurai* (2003).

It is difficult to overestimate the importance of match moving to modern special effects production. The ability to synchronize the performance of virtual and real-world cameras and to layer moving digital objects onto real ones is the key to the invisible integration of live action and digital effects. Film-makers are becoming increasingly reliant on the computer to provide not only the expected animated monsters and spaceships, but also set details such as period buildings, trees and replica historical vehicles.



RENDERING

The last process in the production of computer-generated imagery is 'rendering'. Rendering is a highly complex mathematical operation that conjures a completed, high-quality, 2-D image from the mass of instructional data that is generated during the digital production process. It is one of the most important, and probably least understood, parts of the visual effects pipeline.

When a shot has been finished by an artist, it is launched into a render. During this process, the computer studies every piece of information that it has been given about a scene. As it processes this data, the final image is constructed one pixel at a time. Examining a tiny part of the scene, the computer calculates the geometry of the object that the virtual camera is looking at; what animation, texture maps and shaders have been assigned to that object; and the quality and quantity of lights that surround it. Some objects within a scene will only be generated during the rendering process. For example, a character covered in fur will have only a few guide hairs (<228) placed on it during modelling. The renderer will use the information provided by these guide hairs to calculate the position of the thousands of surrounding hairs. These will then be generated as pieces of geometry that will be included in the calculation of a final 2-D image. When every influence has been considered and calculated, the computer produces a number that represents a colour. The colour is assigned to one pixel, which is then placed in a grid. The computer then begins to calculate the colour for the next pixel along. After the process has been repeated several million times, the result is a single computer-generated image which might be recorded directly onto film or composited with other elements to produce a final image.

Rendering is the most time-consuming and processor-intensive task in the production of digital images. Most large special effects facilities have powerful render 'farms' with hundreds, sometimes thousands, of processors dedicated to rendering. At ILM the server room houses over 3,500 AMD processors that work all day to create the company's groundbreaking images. At night the processors in the artists' desktop workstations also become part of the system, resulting in one of the world's most powerful computing networks with over 5,000 processors.

Rendering is often the deciding factor in the quality of the visual effects that appear in a film. Using modern modelling and animation software, digital artists have the potential to create almost any conceivable image. However, it is the time that such images may take to render that often affects how much can be achieved on the budget of a feature film. The amount of information contained in a scene directly affects the time that it takes to render, and much effort is put into finding ways to make shots more efficient to render. One frame of complex computer animation might typically take four hours to render, though some can take many more. With 120 frames in a five-second sequence, such a shot would take one processor 480 hours – 20 days – to render. Using many processors, however, the shot could be completed overnight. Although the power of processors increases yearly, the growing complexity of the scenes that they are expected to handle means that render times are not much faster now than they were a decade ago.

The final version of a shot is usually rendered in several different layers or 'passes'. Each of these will render a different aspect of the scene, such as the specular, diffuse, or ambient lighting (<166), ambient occlusion (<228), Z-depth information (<103), and mattes. The result will be a number of elements that can be composited together to produce the final image at a later stage. By breaking a shot into its constituent parts, far greater versatility is possible during compositing.


Many types of rendering software are available on the market and visual effects facilities tend to employ several types, depending upon the specific requirements of their work. By far the most popular is Pixar's RenderMan. RenderMan was one of the first effective pieces of rendering software and was developed at Lucasfilm in the mid-80s by Dr Edwin Catmull, now the executive vice-president of Pixar. It was this software that allowed breakthroughs in computer-generated imagery for films such as *The Abyss* (1989) and *Jurassic Park* (1993) and which continues to be used for much of the most spectacular feature film work.

RENDERING LIGHT

An important part of the rendering process is calculating the way that any lights placed in a scene will affect the final image. The way that light is absorbed, reflected or refracted within a CG scene can have considerable impact on the appearance and ultimate realism of a shot.

At a basic level, lighting calculations are achieved by combining information about the quality and angle of lights pointing into a scene with information about objects in the scene as defined by their geometric attributes and the texture maps and shaders that are assigned to them.

Until relatively recently the realism of lighting in a scene was largely a result of the way that the virtual lights were set up by the lighting designer. This was because the computer had no efficient, physically accurate way of calculating the properties of light in a digital scene, meaning designers had to 'fake' the way that light works in the real world (<231). A great deal of academic



LEFT: Many elements are typically rendered for each CG object in a shot. Here are some of those rendered by Luma Pictures to produce a shot of a werewolf for *Underworld: Evolution* (2006).

- a: ambient light pass
- b: diffuse light pass
- c: specular light pass
- d: subsurface-scattering pass
- e: colour pass
- f: fur pass
- g: original live action
- h: final composite



research has gone into finding ways of accurately re-creating the physical behaviour of real light within computer-generated environments. The aim of such techniques is to produce natural-looking lighting by simulating the way that all light sources and object surfaces within an environment interact with one another. This is achieved by accounting for all possible combinations of diffuse and specular reflections and transmissions. Such techniques are grouped under the umbrella term 'global illumination'.

Using global illumination methods a CG environment can, in theory, be illuminated using only a few lights to represent the sources of direct illumination that would naturally occur in that scene. The light from these direct sources then bounces around in the environment to create the diffuse (soft) interreflected illumination that most objects in the real world are lit by. An outdoor scene should therefore only need to be lit by the sky, while an indoor scene should only need lights placed in the windows and at any internal sources of illumination – such as ceiling lamps. By rendering such environments using global illumination algorithms, the light from each source bounces around the scene in accordance with the laws of physics (or at least models of those laws) to produce very realistic lighting. In reality this natural-looking light is rarely sufficient for the purposes of film-making and, just like cinematographers filming outside on a sunny day, the lighting designer will usually place additional artificial lights in order to achieve the dramatic look they require.

Global illumination is a very general term for a number of techniques that aim to create natural-looking lighting by applying the physics of light to the rendering process. Different visual effects companies and rendering packages use a combination of these approaches in order to achieve what they consider to be global illumination.

RAY TRACING

One method of producing realistic-looking images that include accurate reflections is to use a rendering method called 'ray tracing'. To determine the colour of a pixel, a ray is traced from the virtual camera to the spot in the scene that is being studied (fig. 11, 238>). According to the quality of the surface that the ray hits, as described by an object's geometry and the texture maps and shaders that have been assigned to it, the ray will either be absorbed, reflected, or, in the case of translucent materials, refracted.


Any light that is reflected will continue its journey bouncing through the scene until it is either absorbed, exits the scene, or reaches a light source. When the ray has finished its journey, the colour of the single pixel that it represents is adjusted according to how the light has been affected during the ray's journey. This process has to be repeated for each of the millions of pixels in a frame of computer animation and is view-dependent, meaning that the lighting calculations need to be recomputed as soon as the camera moves. Ray tracing is actually the reverse of what happens in

ABOVE: Global illumination techniques now make the creation of natural-looking CG environments possible, as this New York street scene from *King Kong* (2005) illustrates. While the foreground buildings, road, traffic and pedestrians were filmed outside, most of the buildings and distant detail are computer-generated.


BELOW: The machine room at The Motion Picture Company is an atmosphere-controlled environment where hundreds of processors work to create spectacular visual effects. The massive banks of servers, processors and hard disk arrays make this feel more like the control room of NASA than a place where movies are made.



FIGURE 11 RAY TRACING



rendered ray-traced image from position of virtual camera



the real world where photons are emitted by a light source and travel through a scene until they reach our eye.

Ray tracing is particularly useful for calculating reflections and shiny surfaces. The technique does not, however, calculate the diffuse interreflection of light between various surfaces and it produces very sharp shadows that tend to look unrealistic in natural environments.

RADIOSITY

Radiosity was one of the first algorithms developed to compute both direct and indirect illumination.

Radiosity works by dividing the surfaces in a scene into thousands of tiny patches or tiles (fig. 12). Each of these tiles, which start life unilluminated and therefore black, looks out into the scene and gathers information about any light source that it can see. It then calculates a colour and brightness for its

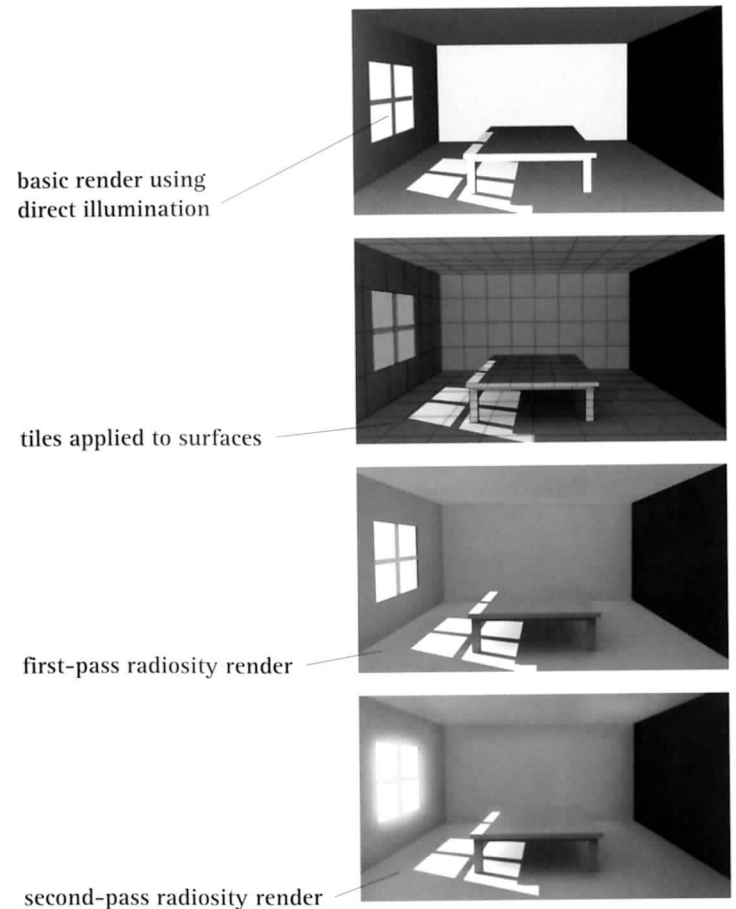
surface. When a tile first looks into a scene everything will be dark except for any sources of direct illumination (lights). The tile will therefore shade itself according to the information coming only from those lights. However, once every tile in a scene has made its initial response to the direct light that they can see, the whole scene will immediately look very different. Now there will be two types of available illumination in the scene: the original direct light sources and the thousands of lit tiles which are themselves now indirect light sources. The lighting calculation is therefore repeated, each tile taking into account the colour and intensity of both the direct illumination (the lights) and indirect illumination (the other tiles) in the scene. This is a recursive process and after each repeated calculation, or pass, the lighting in a scene will become more subtle and realistic. After a number of passes (perhaps several dozen) the changes become increasingly fine and so the process is stopped.

Unlike ray tracing, radiosity calculations are independent of viewer position, so providing that none of the lights or objects within an environment moves, only one, very demanding, calculation is necessary per scene. Once a radiosity calculation is complete the resulting lighting effects are 'baked' into the surfaces in a scene. Wherever that environment is used in a film the same radiosity calculations can be used to light the surfaces during rendering.


Calculating radiosity is an extremely intensive process and is impractical for use in every situation. It is often used for the production of digital matte paintings and environment models (260>).

FIGURE 12 RADIOSITY

After dividing a scene into coloured tiles, several cycles of radiosity are rendered, each creating a more refined and realistic image.



Vital to the realism of scenes containing water and glass, caustic rendering simulates the way that light is refracted and focused by transparent materials.



PHOTON MAPPING

Photon mapping offers another method of calculating the diffuse interreflection of light between surfaces in a computer-generated environment. Using this technique millions of photons (particles of light energy) are fired in random directions from all the lights in a scene. Each photon has a known amount of energy and moves through the scene until it hits an object. The geometry and shaders of that object determine how much of the photon's energy is absorbed, refracted or reflected. They will also determine the angle in which any reflected energy is redirected. The photons continue bouncing from surface to surface within the scene until all their energy is absorbed.

As the photons bounce around in a scene a record of their behaviour is stored as a 'photon map'. This map details the points at which photons hit a surface, their incoming angle and their energy upon arrival.

Photon mapping is a preprocess that is calculated before the main rendering of an image. Once a map has been calculated in a 'first pass' it is stored and then referenced during a later 'second pass' rendering of each frame. The second pass rendering calculates the direct lighting in a scene using standard ray tracing (<237). The indirect light is also calculated using ray tracing but when a ray hits the surface of an object it references the photon map to discover the intensity of light at that point.

As well as producing very natural interreflected diffuse lighting in a scene, photon mapping is very good at calculating a particular lighting phenomenon called 'caustics' (fig. 13). Caustics are formed by light that is reflected or transmitted by a number of specular surfaces before hitting a diffuse surface. Examples of caustics with which we are all familiar are the shimmering patterns of light on the bottom of a swimming pool and the way that light can be focused through a magnifying glass.

Other important lighting effects include subsurface scattering, used to simulate the realistic appearance of translucent skin (<204), and volume rendering, which uses methods similar to ray tracing to calculate the appearance of the large numbers of particles used to form smoke, fire or clouds.

OTHER RENDERING TASKS

In addition to lighting effects, many of the most subtle yet important visual qualities of a scene are generated during the rendering process. These qualities can make the difference between mediocre-looking graphics and successful photorealistic images.

Motion Blur


All fast-moving real-life objects filmed with a movie camera feature motion blur, a streaky effect that results from the subject's movement during exposure. Since computer models do not actually move while they are being photographed (or rather rendered), fast-moving objects can have natural-looking motion blur added to them during the rendering process. The computer, knowing an object's position in previous key frames and in-between frames, is able to calculate how much blur should emanate from an object according to the speed at which it is travelling.

Anti-Aliasing

Anti-aliasing helps to make the 'rough edges' of computer-generated objects look more smooth and natural. A digital image is made up of many thousands of discrete square pixels, each of which can only represent a single colour. In some situations, the picture definition created by these tiny pixels is not good enough to make the edges of an object look naturally smooth, resulting in jagged edges. During rendering, the computer can study the edges of objects and set the characteristics of pixels on boundary edges to be a mixture of the background and foreground qualities, resulting in much smoother edges (fig. 14).

When a shot has been rendered, its frames are assembled in the correct order and the sequence is either viewed on a monitor or transferred to videotape. Once the work has been approved by the film's effects supervisor and director, it is combined with other elements during compositing or recorded onto celluloid film, or to a digital storage device, ready to be edited into the final version of the movie.

FIGURE 14 ANTI-ALIASING



STAR WARS: EPISODE I THE PHANTOM MENACE

Star Wars (1977) and its two sequels, *The Empire Strikes Back* (1980) and *Return of the Jedi* (1983), were among the most innovative and important special effects films ever made. Fans who knew these movies were intended as the middle chapters of a planned nine-part saga waited patiently for more sequels. But for 15 years George Lucas (<39) insisted that effects technology was not yet capable of creating the new worlds he had in mind.

When ILM created dazzling computer-generated dinosaurs for *Jurassic Park* (1993), Lucas realized that technology had finally caught up with his imagination, and he began work on the very first chapter of the story: *Star Wars: Episode I The Phantom Menace* (1999). In terms of scale and ambition, this was to be the greatest visual effects film yet made. Of the 2,000 shots in the film, some 1,900 would be digitally generated or enhanced (a major effects film might normally contain 200 effects shots).

Where previous *Star Wars* films had always been heavily populated by a mixture of lovable and loathsome creatures – usually produced using make-up, animatronics or stop-motion – many of the stars of *The Phantom Menace* were entirely digital, walking, talking and interacting with all the freedom of human actors. The most complex digital character was the clumsy alien Jar Jar Binks. With a face capable of over 400 expressions, Jar Jar was the most emotive and expressive digital character yet created for a movie. Other notable digital characters included the hovering, trunk-nosed junk dealer, Watto, and the toad-like Gungan leader, Boss Nass. Sophisticated cloth-simulation software was developed to create the convincing costumes worn by these and other digital stars. Despite the film's awesome digital menagerie, many characters were still realized using more traditional techniques. These included Yoda, who, two decades after his screen debut, was still simply a hand puppet operated by ex-Muppet performer Frank Oz.

More than any film before, *The Phantom Menace* explored the potential of the 'digital backlot' (262>), filming actors and minimal sets against a blue screen and adding digital or miniature scenery during post-production. Locations such as the planet-city of Coruscant were created almost entirely within the computer and then composited with the live action. Other locations, such as the ornate city of Theed, were traditional models filmed with motion-control cameras and combined with digital extensions. Some methods used to bring locations to life were charmingly old-fashioned, however. The magnificent waterfalls that surround Theed were created by filming dry table salt as it was poured over black velvet – a traditional effects trick.

Star Wars films have always relied on miniatures and models and *The Phantom Menace* was no exception. The film was one of ILM's biggest miniature assignments ever, with hundreds of spacecraft and locations being constructed under the watchful eyes of model supervisor Steve Gawley and chief model-maker Lorne Peterson (<153). While many flying spaceships were digitally animated, those that needed to be destroyed were still blown up using miniature pyrotechnics and filmed at high speed with motion-control cameras.

Some of the film's most interesting effects were invisible. During editing, Lucas was able to radically alter images and even create entirely new ones with remarkable freedom. Many original scenes were digitally manipulated, allowing shots to be assembled from several different takes, choosing each actor's best performance. Some dialogue was rewritten after filming, so actors' mouths were filmed speaking new lines and composited over original footage. Entirely new scenes were created by selecting and manipulating images of performers taken from various scenes and combining them with new backgrounds.

With *The Phantom Menace* and its even more spectacular sequels, *Attack of the Clones* (2002) and *Revenge of the Sith* (2005), George Lucas achieved what he had always dreamt of – total creative freedom without any of the physical and optical constraints that have hampered film-makers for over a century. These were perhaps a glimpse of the way that all films will be made in cinema's second century.



First published in Great Britain 2006 by Aurum Press Ltd
25 Bedford Avenue, London WC1B 3AT
www.aurumpress.co.uk

Copyright © 2006 by Richard Rickitt

The right of Richard Rickitt to be identified as the author of this work has been asserted in accordance with the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this book may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without permission in writing from Aurum Press Ltd.

The author and publishers have made every effort to contact the copyright holders for the illustrations used in this book. The publishers should be notified of any omission of credit.

A catalogue record for this book is available from the British Library.

ISBN-10 1 84513 130 4

ISBN-13 978 1 84513 130 2

10 9 8 7 6 5 4 3 2 1

2010 2009 2008 2007 2006

Design by Ashley Western
Printed in China

FRONTISPIECE: *King Kong* (2005), Weta's extraordinary creation.

BELOW: Special effects have played a prominent part in many movies during cinema's first century: *Santa Claus* (1898), *T2: 3-D - Battle across Time* (1996), *Manslaughter* (1922), *When Worlds Collide* (1951), *The Living Daylights* (1987), *The Godfather* (1972).

