



# What are you in the mood for?

## An introduction to the potential of emotion-based recommendations on streaming platforms

Elina Yli-Ojanperä

#### Contents

Preface	2
Pressing play to regulate emotions	3
Tagging moods	4
The limits of tagging	7
Reactions by emoticons and written words	9
Focus groups	10
Cognitive AI and video analysis	11
What are the ethical risks of analysing viewers emotions, and mood based recommendations?	13
Benefits of emotion based recommendation for creators, publishers and the audience	15
What next?	16
Conclusion	18



Elina Yli-Ojanperä is a Publishing executive at the Finnish public media service company Yle. She specializes in working with digital content and platforms to provide an objective view from the customer's perspective. Her expertise includes supporting digital content development, creating publishing plans, and understanding digital and streaming platform features and data analysis. She has a strong interest in series, digitalization, emotions, and the integration of artificial intelligence with human involvement. She follows the emotional impact from storytelling to publishing and offering content. Contact: elina.yli-ojanpera@yle.fi

This paper was produced as part of the Polis Newsroom Fellowship scheme supported by Yle. Polis is the international journalism think-tank at the London School of Economics and Political Science. This project was overseen by Polis director, Professor Charlie Beckett. The views in this paper reflect those of the author and are not necessarily those of Yle, Polis or LSE.

#### Preface

Human behaviour is very much driven by emotion. So it follows that emotions also influence our media consumption - such as in how we decide what we want to watch on video streaming platforms. There is currently a great deal of talk about sentiment, emotions, feelings and mood when it comes to analysing and curating streaming platform video content. Understanding the emotional effect of content supports improvements in both content recommendation and engagement as well as content development.

Viewers of streaming services often use their feelings to describe the content they find most relevant to their personal preferences and ask for recommendations in viewing groups for new things to watch, based on the emotional reaction caused by a series, movie, or other content they have just watched.

Viewers will also avoid some content because they will react to it in a negative way.. If producers

know how to analyse the content they publish more widely by looking at its emotional effects, they can offer more relevant recommendations and instead of spending time browsing a service, viewers can concentrate instead on watching the actual content.

This report provides an introduction to how metadata and Artificial Intelligence (AI) could be used to recommend content on a video streaming platform, based on the mood and feel it gives to the audience. It will also suggest ideas for the future in this field of emotional recommendations, using AI with humans in the loop. In addition to its opportunities and benefits, an emotion-based recommendation also raises considerations about ethics and risks. This article is based on expert interviews conducted in October 2022 and follow-up interviews in spring/summer 2023.



# Pressing play to regulate emotions

There are various reasons to turn to a favourite streaming platform: to be entertained, be informed and also in some cases it is a way to 'regulate' our emotions. Music<sup>1</sup> has long been understood as a resource to manage emotions. In the music industry, emotions have been widely understood as a part of the creative artistic process and as a consumer need. So recommendations based on emotions are common on music streaming platforms like Spotify<sup>2</sup>. Listening to music is a way to alter mood - for example, to relieve stress. Subscribers are given playlist recommendations for "Fun Run" or for moments when "Life Sucks"

Video streaming platforms are now recognising they need to understand audience emotions when producing and recommending content.. Films and series produce strong emotions and, Netflix, for example, is using emotional analytics<sup>3</sup> to improve customer experience.

News and other media companies are also aiming to produce content that interests and responds to audience needs, and are seeking more intuitive ways to recommend content that has already been published and keep their catalogues in active use.

So how can we use audience data, metadata and AI technologies to understand audience sentiment and to curate video content for consumers to reflect their emotions?

<sup>1</sup> Wadley, G., Krause, A.E., Liang, J., Wang, Z., & Leong, T.W. (2019). Use of music streaming platforms for emotion regulation by international students. *Proceedings of the 31st Australian Conference on Human-Computer-Interaction*. https://dl.acm.org/doi/10.1145/3369457.3369490

<sup>2</sup> BBC, Mark Savage (2021). Spotify wants to suggest songs based on your emotions https://www.bbc.com/news/entertainment-arts-55839655

<sup>3</sup> TechTarget, George Lawton (2019). How Netflix uses emotional analytics to improve CX https://www.techtarget.com/searchcustomerexperience/feature/How-Netflix-uses-emotional-analytics-to-improve-CX



#### **Tagging moods**

With tagging, video content is labelled with keywords describing types ofgenres, actors, and themes for example. There is also tagging for the mood of the content.

The BBC<sup>4</sup> uses tags to describe editorial tone such as critical, funny, and dark and also audience motivation (inspire me, update me)<sup>5</sup>. At the BBC tagging is also widely used in online articles meaning readers can click on a subject tag to find other articles tagged with the same subject keywords.

Editorial tone tags are also used at the BBC as a tool to build recommendational content groups in the music section on the BBC Sounds<sup>6</sup> app: Chill, Feel Good Tunes, and Focus for example. These editorial tone tags group music programmes and curated music lists together. However these tags do not show up in an individual music programme description or in metadata.

4 Tatjana Mladenovic, Lead Data Manager of BBC Archives Technology & Services. Interview 10/2022 Stuart Jennings, BBC lead data scientist. Interview 11/2022 Jeremy Tarling, BBC (2020). How the BBC is approaching metadata quality measurement https://www.iskouk.org/resources/Documents/Events/Dloads/How%20the%20BBC%20is%20approaching%20 metadata%20quality%20measurement.pdf

5 https://www.ftstrategies.com/en-gb/insights/user-needs-a-way-for-newsrooms-to-do-more-with-less/

6 BBC Sounds https://www.bbc.co.uk/sounds



Tagging is widely used in articles produced by the Finnish public service broadcasting company Yle as well as on streaming platform content in Yle Areena<sup>7</sup>.

In Yle Areena atmosphere tags are also used to build and present recommendation strips in the same way as editorial tone tags are used in BBC Sounds. Collecting series and programmes on a strip, based on mood tags, is a way to recommend similar types of content and give some other perspective besides genre type.

The "Nordic excitement" (Pohjoismaista jännitystä) strip collects together series with the atmosphere tag "exciting" and Nordic country of origin for example.



In February 2023 selected tags describing genre and atmosphere were brought to Yle Areena's user interface. The Screenshot below shows an example of *The Great Pottery Throw Down* (In Finnish: Suuri keramiikkakisa) series site<sup>8</sup>. Below the title there are two genre tags (reality = reality, kisailu = competition) and two atmosphere tags (lämminhenkinen = warm-hearted, hyväntuulinen = good-natured). At the moment these tags are describing words without interaction. But in the future they might be a way for users to categorise more similar content.



Besides tagged keywords, metadata such as a video title and a description can be analysed to understand the topic, genre, and emotions associated with a video. This can be used to recommend similar videos to users based on their viewing history and preferences.

8 Yle Areena: Suuri keramiikkakisa (The Great Pottery Throw Down) series page https://areena.yle.fi/1-50466843

## The limits of tagging

Mood recommendation is currently often based on manually added mood keywords.

At best, the tags are very descriptive and they help summarise the style and feel of the content. But they may be added by people who have not watched the content. Instead, they might only infer or guess the mood based on factors such as genre or descriptive text.

The tagging process is often the responsibility of those who know the content best because they have developed and produced it. But sometimes the people who know the content best aren't the most objective when thinking about recommendations. They may want their content to be visible on different parts of the streaming platform, so they add more tags without considering the promises these tags make to the audience.

Traditional broadcasters often have plenty of metadata describing content as part of the organisation's structure, for example for budgeting purposes. But that isn't the most useful or inspiring way to describe content for the audience.

In order to help the tagging process there are solutions for metadata analysis which can add or suggest optimal tags. For example the BBC has a tool that gives suggestions to people who manually add the keywords. But also these solutions are dependent on the quality of used metadata.

There are also commercial tools<sup>9</sup> on the market that use metadata such as description texts, author information, and genre in sentiment analysis. And for already published content, sentiment analysis based on metadata can be combined with an analysis of audience reactions on social media.



Professional Services Executive and metadata expert Matt Eaton<sup>10</sup> gives various examples of categorising metadata:

Viewers turning off is really a good indicator of mood, and that data can be used to say, 'Okay, they turned off when the music got really loud and it was a very dramatic scene.' Maybe we need to recommend something a bit calmer for the next piece of content. I don't know anyone doing that, but sentiment analysis could be used in that way.

Eaton adds you can use opportunities like the 15-second countdown at the end of an episode to give the viewer a moment to think about what's coming next. You could have metadata about the title that could include the synopsis, Internet Movie Database (IMDb) information, and other sources that can be analysed. Additionally, there are factors like duration and the year of production, which are part of the metadata. These aspects don't require complex machine learning but can serve as valuable data sources. They can contribute to establishing a rhythm and help determine the nature of a movie.

10 Matt Eaton, Professional Services Executive, Metadata expert. Interview 10/2022.





# Reactions by emoticons and written words

After publishing, audience reactions may be collected via streaming platforms and social media. Streaming media companies may encourage viewers to provide feedback through tags or emoticons on social media, or their own streaming platforms. This feedback can then be used to offer personalised recommendations. These methods are similar to giving 'likes' or 'hearts' to content you enjoy, and in some instances, there may be a broader range of emoticons available to express feedback, such as defining the emotions evoked by the content.

But it is important to keep in mind that different age groups interpret emoticons in different ways. For example a smiley face could mean either happiness or sarcasm, while a skull may signify either death or humour depending on the viewer's age and cultural background.

The current approach involves tagging external variables and adding metadata. For instance sentiment analysis can be done on any social media platform that provides large amounts of text data for natural language processing. Al consultant Oguzhan Gencoglu says it is important to note that this kind of analysis can be 'noisy'. It can be difficult to tell whether the sentiment is directed towards the show or something else.

Collecting sentiment data all day and manually annotating it is resource-intensive, so using a rough idea of average sentiments or leveraging machine learning and AI algorithms is a common approach. As a result, incorporating more external metadata into the analysis is becoming more prevalent.

#### **Focus groups**

Focus groups watch content and provide feedback as a person or group, or via videomeet and sometimes via surveys. For example Netflix is testing some of it's content with small U.S. subscriber panels<sup>11</sup> where contributors fill out a survey of the content they've seen. Netflix producers re-developed the film *Don't Look Up* to make it more light-hearted following viewer feedback<sup>12</sup>.

When focus groups meet live or via video, these occasions also serve as an opportunity to gather defining moods, such as by observing facial expressions and physical signals to gain perceptions of the emotional reactions, on top of verbal feedback. These mood-analysing tools are based on actual audience reaction measurements, following their heart-rate, eye-movement, and facial expressions.

11 Variety (5/2022): Netflix Has Been Quietly Screening Movies, TV Shows Months in Advance for Subscriber Feedback
12 The Wall Street Journal (12/2022): Netflix to Let Tens of Thousands of Subscribers Give Early Feedback on Content



### **Cognitive AI and video analysis**

Cognitive AI is the use of Artificial Intelligence to simulate human thought. It can be used to map the emotional quality of video. It can analyse data on colours, stress levels, sound effects, soundtracks and music, scene length, camera angles, environments and objects, emotional structure, and camera cuts.

Swedish company Vionlabs and Contentwise in Italy use this kind of video analysis. Vionlabs has built a metadata platform that can analyse a movie and define its elements, such as genre and mood, but also emotions and the stress-level of each scene. Somebody has chosen the music, somebody has chosen the cinematography – everything is designed to affect our emotions. There is an abundance of rich information in video according to Violabs's founder and creative director Arash Pendari:

We have started to explore the art of storytelling and understanding emotions with AI. Essentially, we train computers to watch movies and generate a fingerprint from them. This fingerprint then allows us to generate truly insightful data.

Al predicts human emotions by analysing data. It represents stress levels and positive and negative emotions. This prediction process involves training the Al with input from humans who label a vast number of video clips, indicating whether they perceive them as stressful, happy, sad, and so on. As the Al receives more labelled clips, it is trained to 'understand' how humans perceive emotions in videos.

Pendari says at Vionlabs they use a large and diverse group of people from various cultures and genders in the process, to provide different perspectives. This allows them to gain a comprehensive understanding of emotional reaction, considering the subtle distinctions that may exist between these groups. Al's primary function is to generalise and gauge what the average viewer might perceive as right or frightening. Naturally, what may be scary to one person may not be to another:

66 Our focus lies in evaluating the intended emotions that the directors aim to evoke, rather than my personal reactions as a viewer. It centres on conveying the emotions portrayed in the actual video, rather than my subjective experience on the other side of the screen. 99



Here's an example of how Vionlabs has sorted emotional data for different styles of comedy:

- High positivity with multiple joy-peaks, combined with limited valleys of emotions with low degree of sadness and drama.
- Medium positivity with several joy-peaks, combined with more valleys of emotions with medium degree of sadness and drama.
- Low positivity with few joy-peaks, combined with multiple valleys of emotions with high degree of fear and intensity.

When it comes to categorising any genre, like comedy, there is a wide spectrum of styles within it. A movie like *Sharknado 2* is categorised as a horror comedy, with a comedic element but also intense and bloody scenes. *Rainman* falls more into the realm of a grown-up comedy, with the structure of a drama but including humorous moments. The movie *Dumb and Dumber* falls into the category of pure comedy, mostly indicating funny scenes and only a couple of negative scenes. Besides mood categories, mood tags and similar recommendations, Violabs's tool also offers solutions e.g. for optimal ad-breaks and generating previews.

The Italian company Contentwise has run experiments in cooperation with Vionlabs and Politecnico di Milano. Paolo Cremonesi, an associate professor of Computer Science Engineering at Politecnico di Milano, presented their study *Extracting mise-en-scene and emotional metadata from video content* at the MDN 2022 online conference (annual meeting point for developers and experts working on Metadata and Artificial Intelligence in the Media) run by the European Broadcasting Union EBU.

The objective of the study was to explore the potential of enhancing the user experience (UX) by combining emotional features, automatically extracted from videos, with traditional metadata within an AI-powered recommendation engine. By integrating emotional information into the existing metadata-based system, they aimed to determine if this fusion would lead to substantial improvements in the overall user experience. The study aimed to assess the impact of incorporating emotional features on the recommendation engine and its ability to provide more tailored and satisfying recommendations to users.

Cremonesi has created a multi-modal content-based recommender system that revolutionises traditional metadata by incorporating automatically extracted emotional descriptors from the visual and audio components of videos. This innovative approach aims to enhance the quality of recommendations provided to users.

Cremonesi has conducted user studies to evaluate the effectiveness of recommendations based on emotional descriptors in comparison to metadata-based baselines. In their experiments recommendation accuracy rose by up to 80–100% and the archive usage increased significantly.<sup>13</sup>

<sup>13</sup> Paolo Cremonesi: Extracting mise-en-scene and emotional metadata from video content. MDN 2022 Online Conference. EBU presentations.

 $<sup>\</sup>underline{https://tech.ebu.ch/publications/extracting-mise-en-scene-and-emotional-metadata-from-video-content}$ 

# What are the ethical risks of analysing viewer's emotions, and mood based recommendations?

- Being rewarded for negative emotions, leading to more negative content and feelings.
- Rewarding anger and only strong emotional feelings. Simplifying and leading to populism.
- Freedom for feelings: Are people willing to be followed and targeted because of their emotions? Who has access to this emotional data? Concerns regarding privacy and intrusion.
- Recommending content that is some way similar but the connection is tasteless for example food and famine. Not understanding the context for example the relationship between two people who are kissing.
- Over-reliance on Al.

Andrew McStay, Professor of Technology at Bangor University points out that recommendation systems prioritise content that generates higher engagement, even if it promotes negative sentiments. It is important to question whether this approach is beneficial or desirable.<sup>14</sup> Nic Newman, Senior Research Associate at the Reuters Institute also sees there is a risk people are pushed into negative cycles:

We've done a lot of work in the news on avoidance and the engagement of algorithms that push you into more and more negative cycles of harm, and we want to get away from that. The personalisation of news is very complicated, and it's not like music. So fundamentally, I think we have to go back to the audience's needs and then think about what different kinds of recommendations to deliver for different circumstances at different times for different people.

Some studies<sup>16</sup> conducted during the Covid-19 lockdowns suggest that binge-watching can provoke anxiety. We watch to relax but some forms of extended consumption can have the reverse effect. So how can platforms be more aware of users' emotional state?

<sup>14</sup> Andrew McStay, Profesor of Technology Bangor University. Interview 10/2022

<sup>15</sup> Nic Newman, Senior Research Associate at Reuters Institute. Interview 10/2022.

<sup>16</sup> Raza SH, Yousaf M, Sohail F, Munawar R, Ogadimma EC, Siang JMLD. Investigating Binge-Watching Adverse Mental Health Outcomes During Covid-19 Pandemic: Moderating Role of Screen Time for Web Series Using Online Streaming. Psychol Res Behav Manag. 2021 Oct 8

Pendari from Vionlabs thinks that in the future users will be able to ask their voice assistant for movie recommendations based on the desired mood. However, there are concerns regarding privacy and intrusion:

66 Asking users directly about their current mood can be seen as intrusive, so it's important for voice assistants to offer more subtle ways to gauge user preferences, such as suggesting different moods to choose from. The future of voice assistants and their ability to analyse user voice and mood data will depend on factors like GDPR regulations and user privacy preferences. 99<sup>17</sup>

Lead Data Manager Tatjana Mladenovic, from the BBC Archives Technology & Services, warns of the risks in recommending content that appears similar but where, for example, the connection is tasteless. When it comes to subject access, topics such as food and starvation can converge. Mladenovic also points out the importance of audience age:

**66** It's crucial to use this information responsibly and acknowledge that different age groups have varying preferences for what is friendly and emotionally appropriate. For instance, content designed for 5-year-olds should differ from that tailored for adults aged 44. **99**<sup>18</sup>

Metadata expert Matt Eaton thinks the biggest challenge today is that complex systems often lack an understanding of context:

**66** For instance, when they see two people kissing, they may automatically assume it's a romantic gesture, but it could actually be a display of affection between siblings or close friends. To address this, it's important for these systems to become more sophisticated in comprehending the broader context of the content. Instead of simply focusing on isolated actions like a kiss, they should aim to grasp the entire narrative to avoid drawing incorrect conclusions. Therefore, moving forward, enhancing the ability to understand context is a crucial aspect that we should strive for in the development of these systems.

<sup>17</sup> Arash Pendari, Founder and creative director of Violabs. Interview 10/2022

<sup>18</sup> Lead Data Manager Tatjana Mladenovic from BBC Archives Technology & Services. Interview 10/2022.

### Benefits of emotion based recommendation for creators, publishers and the audience

An emotion-based recommendation model could provide a more fluent way for traditional broadcasters to present their content, rather than recommending content based on the traditional organisational hierarchy of genres and departments.

Mike Leverington, Director of Data Experimentation at ITV, says that there is potential to track emotion-based data and use it to re-think how different productions are linked together and what kind of content is provided through recommendations. For example, it could be used as a tool for curating programmes for special occasions such as Christmas and for feel-good moments when different aged people spend time together.<sup>19</sup>

Pendari points out how mood can be key when we share our watching experience. It becomes easier for us to make viewing decisions based on the shared mood rather than a specific genre or theme:

**66** For instance, if I'm watching with my girlfriend or a companion, we find it simpler to determine the desired mood we want to experience together. Whether it's a lighthearted comedy or a romantic movie, we prioritise the emotional atmosphere we wish to cultivate and enjoy as a shared experience.

With better recommendations, people can concentrate on viewing content rather than spend time browsing and trying to find something new. Users are more likely to move out of genre bubbles when handed more relevant search results, meaning they are more likely to use the full archive, which in turn promotes long-term engagement.

Combining emotional feedback with other data sources allows platforms to improve the overall content development process. Understanding audience emotions and preferences can guide the creation of new content that resonates with viewers and the desired emotional responses.

<sup>19</sup> Director of Data Experimentation at IT Mike Leverington. Interview 10/2022.

#### What next?

- Consumers spend more time with the actual content instead of spending time browsing and trying to find something new. Surveys conducted in the U.S. and the UK show it is not uncommon to spend tens of hours per year on streaming platforms browsing what to watch next.
- Following emotional trends and setting these as specifications for ordering and creating content.
- A working tool for content adaptation: 'same scene with a more uplifting tone.'
- Users remixing content with sentiment tone. Remixing culture around sentiment. Making remixes of popular films or scenes and catching the emotional tone.
- Al with humans in the loop. Fine-tuning emotional goals and feelings by the content consumed.
- Mixed metaverse with people showing and labelling emotional expressions.

Generative AI tools such as ChatGPT and the AI image creator Midjourney offer new possibilities for recommendation and content development. AI Expert Oguzhan Gencoglu says this is a 'game-changer':

Imagine, people will be able to write a movie plot and create an entire movie! You can even modify existing movies, like changing the tone or adding brighter elements. It's not just passive consumption anymore; it opens up a whole new level of creativity and engagement from users. Both publishers and users will benefit from these Al-generated models that can create something out of nothing.

He says that this technology will develop rapidly:

**66** Users will have more control to modify content, text, conversations, and even create movies based on their preferences. There will be thousands of content options for users to modify and personalise. It will be an ecosystem of remix culture, more automated and with tools to help unleash your creativity. And for those who just want to consume, they'll have more options and control.



Voice search, which heavily relies on emotional and contextual data, poses difficulties in content discovery. However Arash Pendari from Vionlabs thinks that in the next five to ten years, as voice control becomes a standard feature, navigating through a user experience solely based on voice commands will become much easier. For example, users will be able to ask their voice assistant for movie recommendations based on the desired mood, making it a crucial aspect of the future of voice assistants.

Pendari also thinks that in the next ten years, AI will continue to advance rapidly, leading to significant changes in the landscape of content production:

66 This democratisation of content production means that we can expect an explosion of movies and TV series originating from sources outside of Hollywood. Al will play a crucial role in empowering creative individuals worldwide, ushering in a new era of diverse and innovative storytelling.

Currently, there's much hype surrounding mixed reality or the 'metaverse', and while its adoption and impact are still debatable, major players like Apple and others are actively developing mixed reality products with emotion-based features. Andrew McStay from Bangor University says that within the next 10 years, mixed reality will see significant advances and an increase in emotion-based services:

**66** For example, in a mixed reality environment, your emotions, such as smiling, could be rendered and reflected in your virtual avatar. This means that as we communicate through screens, our avatars will mimic our real-life expressions, creating a more engaging and natural interaction. **99** 

Striking the right balance between enabling interactive experiences and avoiding intrusive labelling will be crucial for the successful development of mixed reality.

Nic Newman from the Reuters Institute says that we already have the components for the development of recommendations through emotions:

I don't think there's a single new thing that's going to come along and revolutionise everything. I believe all the ingredients for recommendations are already here. We just need to put them together in the right way that benefits us and our society. Many people are currently unhappy with the algorithms. We've spent the last ten years becoming somewhat enamoured with algorithms that constantly provided us with relevant and emotionally satisfying content. However, we are now realising the dangers and downsides of that approach. So, I don't think the answer lies in evoking even stronger emotions. Instead, it may involve incorporating more human judgement and moving towards a more sophisticated and balanced way of assembling recommendations.



#### **Conclusion:**

Emotions are core to our identity as humans but we often find them difficult to explain or understand. How viewers of content behave is influenced by their emotions but it is often complex. So recommendations can only suggest - , they will never be perfect. But where recommendations do work, especially with interactivity as part of the process, they can help the producer understand the emotional impact of their content and it can help the user save time and optimise their viewing experience. Recommendations can help in the development of new content and allow the user to find material that suits their emotional needs.

#### Elina Yli-Ojanperä

Publishing Executive Yle Creative Content and Media This paper represents the views of the author and not necessarily those of Polis or the LSE

Thanks to the following who gave interviews for this paper: Matt Eaton, Professional Services Executive, Metadata expert Oguzhan Gencoglu, AI Expert Mike Leverigton, Director of Data Experimentation, ITV Tatjana Mladenovic, Lead Data Manager, BBC Archives Technology & Services Arash Pendari, Founder and creative director of Vionlabs Andrew McStay, Professor of Technology Bangor University Nic Newman, Senior Research Associate at Reuters Institute Stuart Jennings, Lead Data Scientist, BBC Kim Viljanen, Metadata Expert, Yle Mikael Hindsberg, Business Designer, Yle



#### Ise.ac.uk/media-and-communications

#### f 🔘

Sign up for the Polis Newsletter: http://eepurl.com/t9QYv

#### Polis

Department of Media and Communications The London School of Economics and Political Science Houghton Street London WC2A 2AE

Email: polis@lse.ac.uk

Telephone: +44 (0)20 7107 5176

blogs.lse.ac.uk/polis/

The London School of Economics and Political Science is a School of the University of London. It is a charity and is incorporated in England as a company limited by guarantee under the Companies Acts (Reg no 70527).

The School seeks to ensure that people are treated equitably, regardless of age, disability, race, nationality, ethnic or national origin, gender, religion, sexual orientation or personal circumstances.

Design: LSE Design Unit (info.lse.ac.uk/staff/divisions/ communications-division/design-unit)