IBC2022

PREDICTING ENERGY USAGE OF HIGH DYNAMIC RANGE VIDEO ON MOBILE DEVICES

Daniel Ilett[1], Kurt Debattista[1], Mike Nilsson[2], Paul Farrow[2], and Alan Chalmers[1]

[1] University of Warwick, [2] British Telecommunications plc, United Kingdom

ABSTRACT

High-end mobile devices now support displaying video in High Dynamic Range (HDR), delivering a significantly enhanced viewing experience over Standard Dynamic Range (SDR). However, more energy may be required to play HDR, impacting device battery life and reducing overall quality of experience.

We present a new methodology for predicting the real-time energy usage of a mobile device playing video content. 37 video clips were encoded into 12 combinations of different resolution, frame-rate, bit-rate, and dynamic range. An external power monitor was used to measure the voltage and current drawn by the device while playing the content. These measurements were used to train a neural network to predict the energy requirements of playing any clip.

We show that our model can predict the energy usage of videos with RMS error of 4.88%, achieving a substantial improvement over existing methods that use linear regression, symbolic regression, or trust-region optimisation.

INTRODUCTION

Video content producers and platform owners strive to deliver the best possible viewing experience to customers. Steady improvements to hardware have resulted in mobile devices holding the largest share of the video streaming market [1]. One of those improvements is the availability on phones of High Dynamic Range (HDR) video, an imaging technology that provides greater representation of real-world lighting compared to the traditional Low or Standard Dynamic Range (SDR) imaging. Further hardware enhancements, which include greater screen resolutions and faster wireless streaming, outpace the rate of improvements in Lithium-ion battery technology which is now reaching its theoretical limits [2]. Unfortunately, future portable battery technologies which could replace Li-ion are still in development and are not ready for commercial deployment [3].

In this paper, we describe the method we used to accurately measure the energy usage of a mobile device playing video content. These measurements were used to train a neural network to predict the energy usage of a video clip based on its intrinsic properties: the dynamic range, frame rate, resolution, bit-rate, and pixel luminance (brightness) distribution.

Our model can be combined with extra information, such as the remaining video duration and device battery capacity, to create a system which balances the viewing quality with the need to conserve battery. With the pervasive use of mobile devices in modern life, preserving battery capacity is something that many more users are aware of. Unfortunately, video streaming consumes a huge amount of energy due to the use of the screen, which is "the dominant power consumer in battery-operated devices" [4].

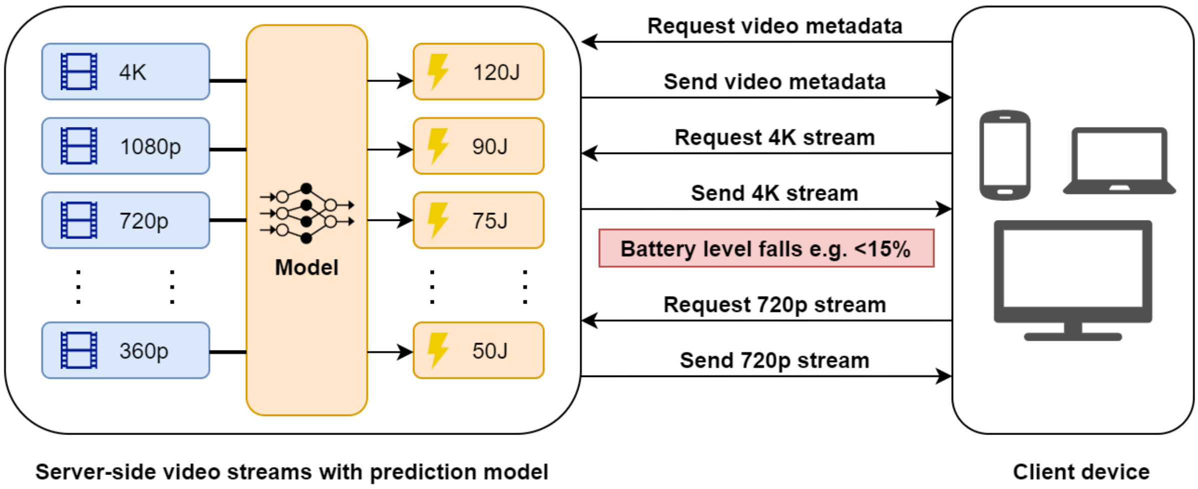

Figure 1 - Diagram of an adaptive system which incorporates energy predictions.

The key contribution of this paper is a neural network-based method to predict the energy usage of a mobile device playing video content, relying only on the properties of the video content itself. This model can be seen on the left-hand side of Figure *1*. Previous work in this area has focused solely on SDR content [5]. This paper describes a model that applies to both SDR and HDR content, and which is also more accurate than previous methods.

RELATED WORK

This work spans multiple fields of research, including video streaming, power monitoring on mobile devices and HDR content. Previous research in each field is discussed briefly below.

High Dynamic Range Video

SDR video is only capable of representing around 8 stops of dynamic range, whereas the human visual system can perceive around 14 stops simultaneously [6]. HDR imaging represents a broader range of luminance values than SDR and can exceed 14 stops. HDR video has gradually become more feasible on mobile devices due to several advancements in technology [7], in particular the HDR video compression standards HDR10, HDR10+ and Dolby Vision [8]. Before this, most consumer displays (including those on mobile devices) were not capable of natively displaying HDR content. Instead, a range of tone mapping operators (TMOs) were used to map the luminance values of HDR content to the luminance range that an SDR display is capable of reproducing [9].

Pramanik et al. [10] conducted a state-of-the-art review of smartphone energy usage in four key areas: power modelling, power management, battery development, and battery hazards. None of the studies reviewed in this paper focused on the impact of HDR video content on energy usage.

Video Streaming

The predominant streaming technology for adapting video quality in response to fluctuating bandwidth availability is Adaptive Bit-Rate (ABR) streaming, of which DASH [11] and HLS [12] are two examples. With ABR streaming, a video is encoded into several streams of different bit-rate (and therefore different quality), with each stream being stored as multiple small media segments. When a client device requests a new media segment, the system can dynamically choose a quality level based on the available bandwidth. If the bandwidth drops, the client attempts to download a lower-quality version of the next segment to prevent buffering and other unfavourable behaviour.

Several methods exist to adapt video quality to reduce its energy impact. The system by Kennedy et al. [13], BaSe-AMy, adapts based on the stream duration and packet loss rate, which can reduce battery usage by around 18%. The DEAS system by Ding and Muntean [14], which adapts based on the energy usage characteristics of different device components, shows up to 40% improved energy performance over similar systems. Furthermore, Lee et al. [15] propose a scheme that reduces spatial resolution of the image, which can reduce energy consumption by around 50% compared to conventional adaptation.

It is also possible to adapt based on user preference or engagement. Evidence shows that video buffering leads to a drop in user engagement. Krishnan and Sitaraman [16] showed that viewers experiencing buffering interruptions exceeding 1% of the total video run-time chose to watch 5.02% less of the video than those that did not. Dobrian et al. [17] conclude in their paper that “the bit-rate is especially critical for live (sports) content'”, highlighting the importance of maintaining a balance between high bit-rate and preventing buffering periods when viewing live sports.

Mobile Device Power Monitoring

Three broad strategies for establishing the ground truth energy usage of a process emerge from the literature.

The first, simplest methods utilise the battery percentage indicator on the device's notification bar [18] [19]. This can be useful for comparing the energy usage of long tasks, but it is imprecise for short tasks. Also, as the battery percentage indicator may not accurately represent the underlying battery characteristics [20], it is unsuitable for our work.

The second set of methods involve directly polling the battery's fuel gauge using a software solution to measure real-time voltage and current draw. This data is far more accurate than the battery percentage method but is difficult to reproduce across devices because the fuel gauge cannot always be accessed easily [21]. This approach facilitates real-time battery monitoring, with a small margin of error between it and the true energy usage [22].

The final technique is to use external power measurement hardware, such as sense resistors [23], a multimeter [24], or specialised hardware (e.g. the Monsoon Power Monitor) [25]. These methods are far more accurate than the others, but they typically require physical access to, or replacement of, the battery.

Battery Lifetime Estimation

Carroll and Heiser [23] measured the energy usage of different mobile device components and estimated the total battery life of several usage profiles which use those components in different proportions. Zhao et al. [18] built a context-aware battery lifetime prediction system by using multiple linear regression, with a maximum error of 6%. Building on both works, Nusawat, Adulkasem, and Chantrapornchai [26] demonstrated a data mining approach for battery life prediction, focusing on multilayer perceptron (MLP) and support vector machine (SVM)-based models. A state-of-the-art review of smartphone battery state-of-charge (SoC) prediction papers by Singh et al. [27] identifies the key studies.

METHOD

This section introduces the methodology for our study and presents an overview of each major step. A full overview of the experiment can be seen in Figure 2.

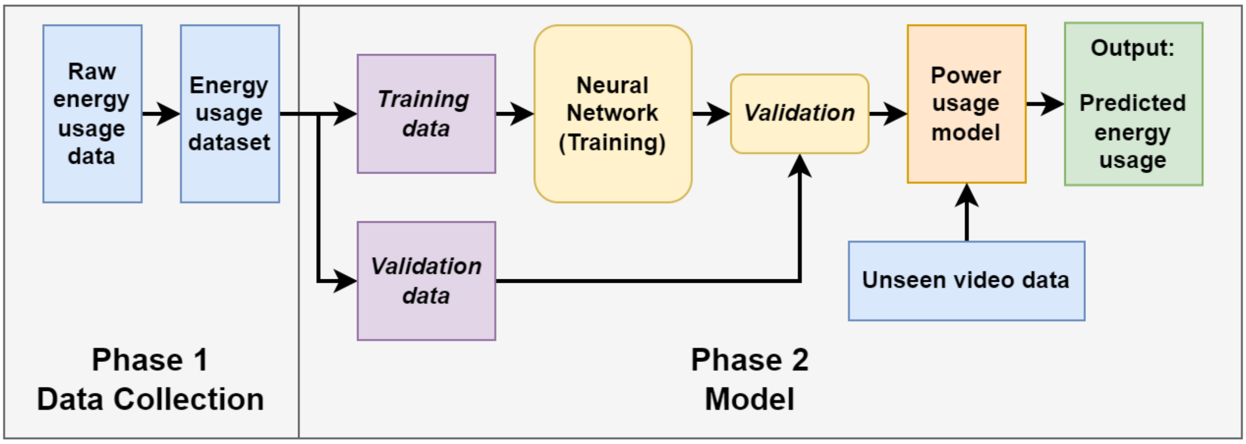

Figure 2 – An overview of the major steps in this study.

Since Lithium-ion batteries have almost reached their theoretical limits, the onus on maintaining reasonable battery life has shifted partially onto software developers, who must balance the technical requirements of their apps, including HDR, with the limited battery life of devices. We may potentially save a significant amount of battery life by adapting video quality based on the energy requirements of each quality level. To this end, the purpose of this study is to create a model which can accurately predict the energy usage of both HDR and SDR video on mobiles. The first step involves building a high-quality dataset containing the real-world energy usage of different types of videos running on a mobile device. This is described in detail in the Data Collection section.

The raw data obtained in this experiment is used to train an artificial neural network. Once the network has been trained, we can use it to predict the energy usage of any video clip, including unseen footage, based only on the properties of the video content itself. This network could theoretically be implemented to predict the energy usage of both on-demand content and live-streamed footage. This phase of the study is described in the Model section.

DATA COLLECTION

Design

We select a set of video clips and encode them to several quality levels. Each clip is then played on an HDR-compatible mobile device and external hardware is used to log the instantaneous voltage and current of the device for the duration of the video. From these values, we can calculate the total energy usage of the device playing the video and build a dataset linking the properties of the video content and the energy usage of that content.

We chose 22 clips provided by BT Sport and 15 clips from the Stuttgart HDR dataset [28]. These clips were chosen due to the contrast in each clip between bright areas, such as sunlight and bright sparks, and low-light shadow areas.

The following subsections set out the variables we are interested in, the equipment we used for the measurements, the control environment, and the measurement procedure.

Variables

Several factors influence the amount of energy used by a mobile device while it is playing the video, including:

- The intrinsic properties of the content itself (bit-rate, frame-rate, resolution, dynamic range, and luminance profile).
- The environment in which the content is played, including other apps on the device.
- The user's preference settings, such as the device's global maximum brightness.

Resolution (pixels)	**Frame rate (fps)**	**Dynamic Range**	**Bit rate (kbit/s)**
1920×1080	50	HDR-10	6500
1920×1080	50	SDR	6500
1280×720	50	HDR-10	5000
1280×720	50	SDR	5000
1280×720	25	HDR-10	2500
1280×720	25	SDR	2500
720×408	25	HDR-10	1800
720×408	25	SDR	1800
720×408	25	HDR-10	1200
720×408	25	SDR	1200
640×360	25	HDR-10	800
640×360	25	SDR	800

Table 1 - Streaming configurations

For this study, we are interested in the effect of changing the bit-rate, frame-rate, resolution, and dynamic range of the content. Video delivery pipelines typically bundle settings into distinct quality profiles, so we chose to encode six quality settings based on the bit-rate, frame-rate, and resolution. Furthermore, we also consider the impact of dynamic range on the energy usage.

'HDR video' is a term that encompasses several formats with a higher dynamic range than SDR content. One of the most widespread of these formats is HDR-10, which itself is a collection of several technologies, including the Perceptual Quantiser (PQ) transfer function [29], a bit-depth of 10 bits per pixel per channel, and the BT.2020 colour space [30]. With that in mind, we encode each clip at each quality level with the H.265 codec in both HDR-10 and SDR using FFmpeg, resulting in 12 quality levels as listed in Table *1*. We use the Main profile for the SDR clips and the Main 10 profile for the HDR clips.

The energy usage of individual pixels on an OLED screen is proportional to the luminance of that pixel [4]. Consequently, the distribution of pixel luminance across the video may greatly influence the energy usage of the video. Our dataset contains a luminance histogram with 10 bins to approximate the distribution of dark and light pixels across the video, along with the mean, median, minimum, maximum, variance, and interquartile range (IQR) of the luminance in the video. An example of one luminance histogram is shown in Figure 3.

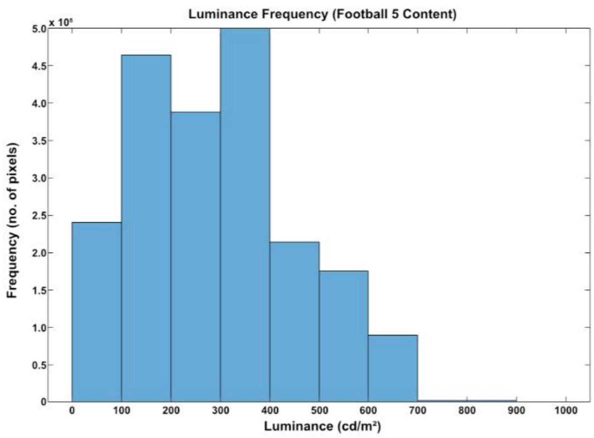

Figure 3 - The luminance histogram and thumbnail for the Football 5 clip.

Equipment

The test device is a Samsung Galaxy S9+, one of the earliest devices to support HDR-10. Its display has a maximum resolution of 2960×1440 pixels and a brightness of 1130nits in high brightness mode [31]. This is far above the peak brightness of SDR devices, which is typically between 100-600nits [7].

A Monsoon High Voltage Power Monitor is used to bypass the device's battery, supply a stable current to the device, and measure the instantaneous voltage and current drawn by the mobile device with a sampling frequency of 500Hz. The power monitor is attached to a computer via USB to log the voltage and current drawn by the device (see Figure 4).

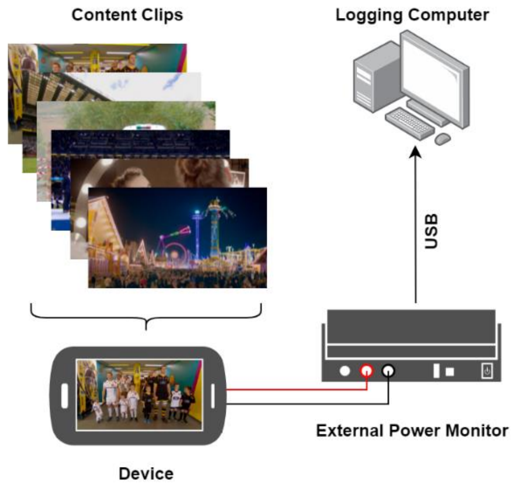

Figure 4 - An overview of the setup for data collection.

Environment

The ambient environment around the device, including temperature and humidity, can impact the energy usage of the device, as can background tasks and other apps running at the same time as the video. The following control environment was used to minimise the effect of other variables on the energy usage.

- Screen brightness was locked to the maximum supported by the device, and adaptive brightness was disabled.
- Screen timeout was locked to the maximum of 10 minutes
- Unrelated hardware features such as Wifi, 4G, Bluetooth, NFC, mobile hotspot, and location services were disabled.
- All user apps were closed, besides any apps directly involved with this experiment.
- The external environment – a temperature of 20°C, humidity of around 40%, and a light level of about 300 lux - was kept constant.

Measurement Procedure

At the start of the data collection phase, the device was connected to the Monsoon Power Monitor and placed in the control environment. The phone was turned on and left alone until the operating system loaded and start-up tasks finished. At this stage, the power monitor recorded a stable current draw of around 100mA by the device.

For each measurement, the phone's default Video Player app was opened, and a video clip was picked at random from the set of clips and played from device memory while the instantaneous voltage and current were logged. The process was repeated until all clips had been played three times. The total energy usage E, in Joules, of the device playing each clip was calculated using the following equation:

$$E_{total} = \frac{1}{500}\sum_{t=1}^{T} V_t \cdot I_t$$

where T is the number of timesteps, V_t and I_t are the instantaneous voltage and current respectively, and $\frac{1}{500}$ is the time, in seconds, between measurements. The mean of the three measurements for each clip was taken to obtain the final energy usage value for that clip. The dataset contained raw energy usage values between 38.9J and 101.5, with a mean of 62.4J and standard deviation of 12.6J.

MODEL

Artificial neural networks (ANNs) [32] [33] are capable of learning non-linear relationships between a set of input and target variables. Once trained, a network can predict a target value from a set of inputs quickly, making it suitable for a system which needs to evaluate the energy usage of streamed video content in real-time. However, due to the nature of neural networks themselves, it will be difficult to determine which variables have the largest overall impact on the energy usage just by looking at the trained model.

This section explores the architecture and performance of the prediction model in detail.

Network Structure

The model, shown in Figure 5, is a feedforward neural network with a single branch of layers with the goal of performing a regression to predict energy usage from the video properties described previously. It comprises an input layer, three hidden layers of size 512, and an output layer. Each layer is itself made up of 'sub-layers' which give them unique behaviour.

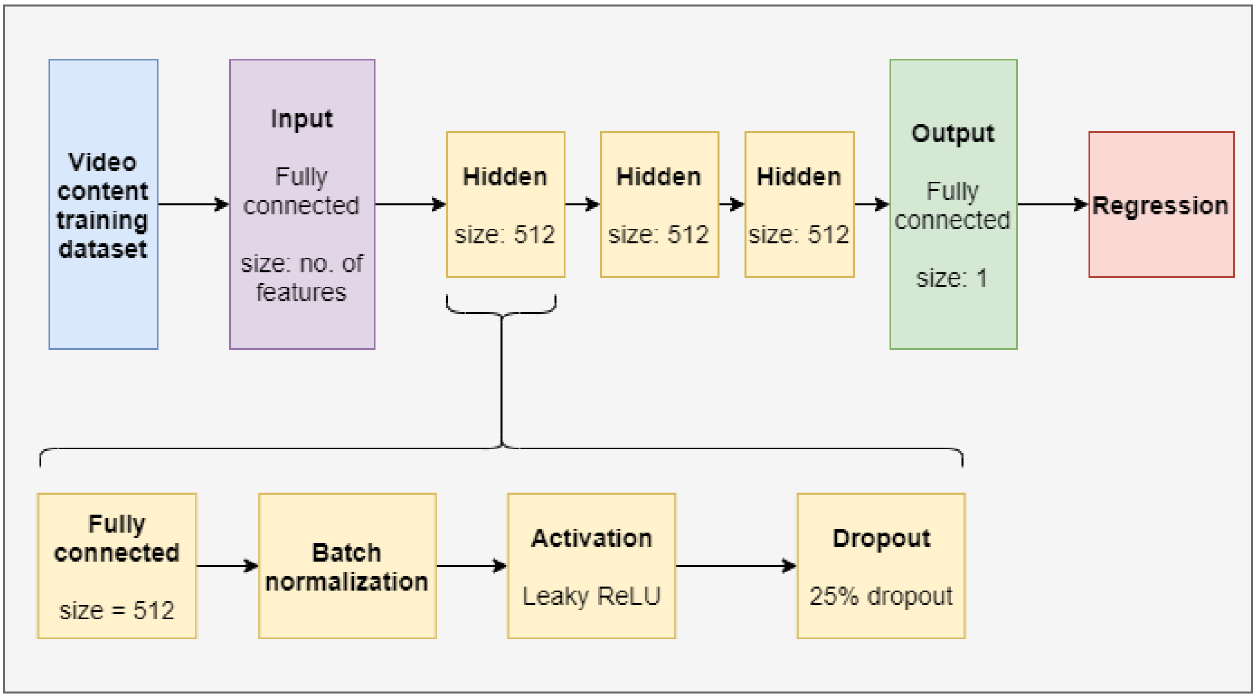

Figure 5 - A diagram showing how each of the layers of the model connect to each other.

The input layer takes each feature as input and applies z-score normalisation so that the network is not sensitive to features with different orders of magnitude. The output layer is a fully connected layer of size one with a regression layer.

The three hidden layers are fully connected layers of size 512. Each one uses batch normalisation to prevent internal covariate shift and 25% dropout to prevent co-dependence between network layers, which could prevent the model from generalising to unseen data.

Activation Function

Each hidden layer uses a Leaky Rectified Linear Unit (ReLU) activation function which introduces non-linearity into the model. With standard ReLU, neurons may have their weights set to zero, which effectively makes them `dead', as seen in this equation:

$$f(x) = f(x) = \begin{cases} x, & \text{if } x > 0 \\ 0, & \text{otherwise} \end{cases}$$

Leaky ReLU solves this problem by introducing a small gradient for values below zero where the ReLU curve usually equals zero, allowing negative activations [34]:

$$f(x) = f(x) = \begin{cases} x, & \text{if } x > 0 \\ 0.01x, & \text{otherwise} \end{cases}$$

Loss Function

The goal of the network is to produce predictions that are as close to the ground truth as possible. A loss function is used to calculate the gap between real and predicted values. For a typical regression problem like this one, mean squared error (MSE) is used as the loss function. The mean squared error is given by:

$$\mathrm{MSE} = \frac{1}{N}\sum_{i=1}^{N}(t_i - y_i)^2$$

where, for this model, there are N observations, t is the vector of predicted energy usage values, and y is the vector of real energy usage values.

Features

The training data comprises the 20 input features listed below, each being a characteristic of an encoded video clip, and the energy usage in Joules as the target feature.

- Pixel count (horizontal and vertical resolution values multiplied together).
- Frame rate, in frames per second (fps).
- Bit-rate, in kilobits per second (kbit/s).
- Bit-depth (10 for HDR-10 footage and 8 for SDR).
- Luminance histogram, where each of the 10 bins is a separate feature.
- Minimum, maximum, mean, median, variance, and IQR of the video luminance, each one being a separate feature.

Training

The model uses 250 epochs with a batch size of 16. This is sufficiently high to ensure the network converges. Learning rate decay is used to avoid the network overshooting (with a high learning rate) or not converging (with a low learning rate). The learning rate for this network starts relatively high at 0.01 and decreases by a factor of 5 every 50 epochs.

The dataset contains 444 observations (37 clips each with 12 settings). We trained with 5-fold cross validation, where the data was divided into five roughly equal partitions and used to train five models. For each model, one partition became the validation dataset and the other four were combined to form the training dataset. The validation partition rotated between models, meaning that each of the five models was trained on around 355 observations and tested with around 89 unseen observations.

Performance

Due to the relatively high number of epochs, each model took a minute or two to train. **Error! Reference source not found.** shows the loss function over the training process, where the steep drop in the first 100 epochs can clearly be seen, followed by the reduced rate of change due to the decreasing learning rate.

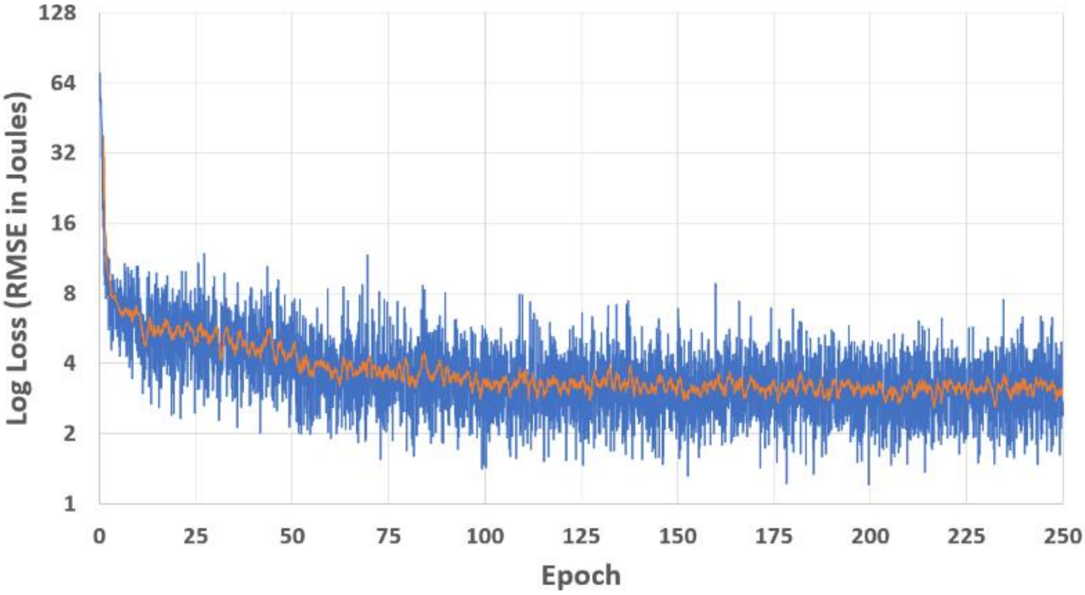

Figure 6 - A plot of the loss function during training. The orange line is a moving average of all training iterations from the last epoch.

For each of the five cross-validated models, we obtained a single root mean squared error (RMSE) value to evaluate the performance of the model, with smaller values indicating better performance.

Table 2 shows a comparison with the earlier prediction methods of linear regression [35], symbolic regression [36], and trust-region optimisation [37].

Method	**Per-model RMSE (J) for the full dataset**					
	1	**2**	**3**	**4**	**5**	**Mean**
Linear Regression [35]	5.9449	8.0751	5.8554	9.6679	6.4167	7.1920
Symbolic Regression [36]	5.7115	6.5024	5.5956	6.1089	5.9332	5.9703
Trust-region Solver [37]	6.5565	6.8244	6.6056	7.1643	6.6969	6.7695
Our Method	**5.4468**	**4.1956**	**4.6342**	**5.7735**	**4.6951**	**4.9494**

Table 2 - Validation error (RMSE) of each type of model.

In terms of mean performance, our neural network-based method performs 31.9% better than linear regression, 26.9% better than trust-region optimisation, and 17.1% better than symbolic regression on the full dataset.

DISCUSSION

The results suggests that the model can accurately predict the energy usage of a section of video content on the test device, whether the device is playing HDR or SDR content.

The energy measurement methodology employed in these experiments can be reproduced across most devices currently on the market, although it is necessary to perform a complicated battery bypass on almost all of them due to the decline in phones with removable batteries. The process we detail in this paper can be employed on any type of video content, including other forms of HDR such as HDR-10+. Furthermore, a model like this is not restrained to just phones; a similar approach could extend to predicting the energy usage on other portable devices, provided the energy usage can be measured and used to calibrate the model.

This model could be incorporated into an existing ABR streaming system, where the choice of video quality for the next video segment would be based not just on the bandwidth requirements, but also the energy requirements and the state of the device's battery.

Limitations

Much of the error in the model's predictions may be attributable to small errors in measuring the energy usage of the mobile device caused by background tasks. Although the video clips were played in a randomised order three times each, it is difficult to entirely remove the impact of these background processes during the measurement phase. Further study could investigate methods for modelling the background task impact separately. This would allow the model to take only the energy usage related to displaying the video content into account and would likely result in even more accurate predictions.

CONCLUSION AND FUTURE WORK

This work provides a methodology for establishing the link between objective video metrics and energy usage on a mobile device. This methodology has been shown to be sufficiently robust in a setting with a relatively small amount of content, and there are potential improvements that can be made to reduce the error rate of the predictions made by the model, as described in the Limitations section. The model may also benefit from additional data points, such as different codecs, which may have an impact on the energy usage of the mobile device while displaying the video content.

The model we presented in this paper predicts energy usage based on video characteristics. To incorporate the findings into existing adaptive streaming systems, future work will be needed to develop a system which can make decisions based on the predictions, as seen in Figure *1*. This includes choosing to adapt to a lower-quality video stream when the device battery runs low, even if the bandwidth available to the device could support higher quality.

However, there may be instances where the energy usage for a piece of content under two different quality profiles is roughly equivalent (for example, a high-resolution SDR stream, compared to a low-resolution HDR stream). In such cases, it is difficult to choose one stream over the other, especially if they also have similar bandwidth requirements. Therefore, further work is needed to build a system which adapts content based on the subjective preference of users in cases where there is a tie between bandwidth or energy usage. This decision procedure could incorporate additional subjective variables such as the ambient viewing environment.

REFERENCES

[1] G. V. Research, 2020. Video Streaming Market Size, Share, Analysis, Industry Report, 2027. https://www.grandviewresearch.com/industry-analysis/video-streaming-market. Accessed: 2020-03-24

[2] J. B. Goodenough and K.-S. Park, 2013. The Li-ion Rechargeable Battery: A Perspective. Journal of the American Chemical Society, vol. 135, no. 4, pp. 1167–1176.

[3] Y. Liu, P. He, and H. Zhou, 2018. Rechargeable Solid-State Li–air and Li–S Batteries: Materials, Construction, and Challenges. Advanced Energy Materials, vol. 8, no. 4, p. 1701602.

[4] D. Shin, Y. Kim, N. Chang, and M. Pedram, 2011. Dynamic Voltage Scaling of OLED Displays. Proceedings of the 48th Design Automation Conference, pp. 53–58.

[5] H. Ahmad, N. Saxena, A. Roy, and P. De, 2018. Battery-Aware Rate Adaptation for Extending Video Streaming Playback Time. Multimedia Tools and Applications, vol. 77, no. 18, pp. 23877–23908.

[6] T. Borer and A. Cotton, 2016. “A Display-Independent High Dynamic Range Television System. SMPTE Motion Imaging Journal, vol. 125, no. 4, pp. 50–56.

[7] F. Banterle, A. Artusi, K. Debattista, and A. Chalmers, 2017. Advanced High Dynamic Range Imaging. AK Peters/CRC Press.

[8] C. Chinnock, 2016. Dolby Vision and HDR10. Insight Media Whitepaper, 2016.

[9] M. Melo, M. Bessa, K. Debattista, and A. Chalmers, 2014. Evaluation of HDR Video Tone Mapping for Mobile Devices. Signal Processing: Image Communication, vol. 29, no. 2, pp. 247–256.

[10] P. K. D. Pramanik, N. Sinhababu, B. Mukherjee, S. Padmanaban, A. Maity, B. K. Upadhyaya, J. B. Holm-Nielsen, and P. Choudhury, 2019. Power Consumption Analysis, Measurement, Management, and Issues. IEEE Access, vol. 7, pp. 182113–182172.

[11] T. Stockhammer, 2011. Dynamic Adaptive Streaming over HTTP – Standards and Design Principles. Proceedings of the second annual ACM conference on Multimedia systems, pp. 133–144.

[12] Apple, 2020. HTTP Live Streaming. https://developer.apple.com/streaming/. Accessed: 2020-03-26.

[13] M. Kennedy, H. Venkataraman, and G.-M. Muntean, 2010. Battery and Stream-Aware Adaptive Multimedia Delivery for Wireless Devices. IEEE Local Computer Network Conference, pp. 843–846, IEEE.

[14] R. Ding and G.-M. Muntean, 2013. Device Characteristics-based Differentiated Energy-Efficient Adaptive Solution for Video Delivery over Heterogeneous Wireless Networks. IEEE Wireless Communications and Networking Conference (WCNC), pp. 4588–4593, IEEE.

[15] H. Lee, Y. Lee, J. Lee, D. Lee, and H. Shin, 2009. Design of a Mobile Video Streaming System using Adaptive Spatial Resolution Control. IEEE Transactions on Consumer Electronics, vol. 55, no. 3, pp. 1682–1689.

[16] S. S. Krishnan and R. K. Sitaraman, 2013. Video Stream Quality Impacts Viewer Behavior: Inferring Causality using Quasi-experimental Designs. IEEE/ACM Transactions on Networking, vol. 21, no. 6, pp. 2001–2014.

[17] F. Dobrian, V. Sekar, A. Awan, I. Stoica, D. Joseph, A. Ganjam, J. Zhan, and H. Zhang, 2011. Understanding the Impact of Video Quality on User Engagement. ACM SIGCOMM Computer Communication Review, vol. 41, no. 4, pp. 362–373.

[18] X. Zhao, Y. Guo, Q. Feng, and X. Chen, 2011. A System Context-aware Approach for Battery Lifetime Prediction in Smart Phones. Proceedings of the 2011 ACM Symposium on Applied Computing, pp. 641–646, ACM.

[19] G. Kalic, I. Bojic, and M. Kusek, 2012. Energy Consumption in Android Phones when Using Wireless Communication Technologies. 2012 Proceedings of the 35th International Convention MIPRO, pp. 754–759, IEEE.

[20] A. Rahmati, A. Qian, and L. Zhong, 2007. Understanding Human-Battery Interaction on Mobile Phones. Proceedings of the 9th International Conference on Human Computer Interaction with Mobile Devices and Services, pp. 265–272.

[21] M. A. Bokhari, B. R. Bruce, B. Alexander, and M. Wagner, 2017. Deep Parameter Optimisation on Android Smartphones for Energy Minimisation. Proceedings of the Genetic and Evolutionary Computation Conference Companion, pp. 1501–1508, ACM.

[22] G. Avvari, B. Pattipati, B. Balasingam, K. Pattipati, and Y. Bar-Shalom, 2015. Experimental Set-up and Procedures to Test and Validate Battery Fuel Gauge Algorithms. Applied energy, vol. 160, pp. 404–418.

[23] A. Carroll, G. Heiser, et al., 2010. An Analysis of Power Consumption in a Smartphone. USENIX annual technical conference, vol. 14, pp. 21–21, Boston, MA.

[24] L. Corral, A. B. Georgiev, A. Sillitti, and G. Succi, 2013. A Method for Characterizing Energy Consumption in Android Smartphones. 2013 2nd International Workshop on Green and Sustainable Software (GREENS), pp. 38–45, IEEE.

[25] J. Chung, H. Jung, J. Koo, Y. Kim, and U.-M. Kim, 2017. A Development of Power Consumption Measurement System for Android Smartphones. Proceedings of the 11th International Conference on Ubiquitous Information Management and Communication, pp. 1–4.

[26] P. Nusawat, S. Adulkasem, and C. Chantrapornchai, 2014. Battery Discharge Rate Prediction Model for Mobile Phone Using Data Mining. 2014 6th International Conference on Knowledge and Smart Technology (KST), pp. 69–74, IEEE.

[27] M. Singh, J. Trivedi, P. Maan, and J. Goyal, 2020. Smartphone Battery State-of-Charge (SoC) Estimation and Battery Lifetime Prediction. 2020 10th International Conference on Cloud Computing, Data Science & Engineering (Confluence), pp. 94–101, IEEE.

[28] J. Froehlich, S. Grandinetti, B. Eberhardt, S. Walter, A. Schilling, and H. Brendel, 2014. Creating Cinematic Wide Gamut HDR-Video for the Evaluation of Tone Mapping Operators and HDR-Displays. Digital Photography X, vol. 9023, pp. 279–288, SPIE.

[29] SMPTE Standard, 2014. High Dynamic Range Electro-Optical Transfer Function of Mastering Reference Displays. SMPTE Standard 2084, pp. 1–14.

[30] M. Sugawara, S.-Y. Choi, and D. Wood, 2014. Ultra-High-Definition Television (Rec. ITU-R BT. 2020): A Generational Leap in the Evolution of Television. IEEE Signal Processing Magazine, vol. 31, no. 3, pp. 170–174.

[31] R. M. Soneira, 2018. Galaxy S9 OLED Display Technology Shoot-out. http://www.displaymate.com/Galaxy_S9_ShootOut_1s.htm. Accessed: 2019-11-20.

[32] W. S. McCulloch and W. Pitts, 1943. A Logical Calculus of the Ideas Immanent in Nervous Activity. The Bulletin of Mathematical Biophysics, vol. 5, no. 4, pp. 115–133.

[33] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, 1985. Learning Internal Representations by Error Propagation. Tech. rep., California Univ San Diego La Jolla Inst for Cognitive Science.

[34] A. L. Maas, A. Y. Hannun, A. Y. Ng, et al., 2013. Rectifier Nonlinearities Improve Neural Network Acoustic Models. Proc. ICML, vol. 30, p. 3, Citeseer.

[35] L. Zhang, B. Tiwana, Z. Qian, Z. Wang, R. P. Dick, Z. M. Mao, and L. Yang, 2010. Accurate Online Power Estimation and Automatic Battery Behavior Based Power Model Generation for Smartphones. Proceedings of the 8th IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis, pp. 105–114.

[36] E. Rattagan, Y.-D. Lin, Y.-C. Lai, E. T.-H. Chu, and K. C.-J. Lin, 2018. Clustering and Symbolic Regression for Power Consumption Estimation on Smartphone Hardware Subsystems. IEEE Transactions on Sustainable Computing, vol. 3, no. 4, pp. 306–317.

[37] C. Herglotz, S. Coulombe, C. Vazquez, A. Vakili, A. Kaup, and J.-C. Grenier, 2020. Power Modeling for Video Streaming Applications on Mobile Devices. IEEE Access, vol. 8, pp. 70234–70244.