



## **REALISING ADDITIONAL VALUE FROM LINEAR CONTENT USING METADATA AND AUTOMATION.**

Michael Armstrong, Miles Bernie, Dave Bevan and Andy Brown

BBC Research & Development, UK

### **ABSTRACT**

This paper describes a collaboration between BBC News and BBC Research & Development which is exploring ways of using the BBC's linear programme output and associated production metadata, along with artificial intelligence and machine learning techniques, to segment the programmes into individual stories. These segments can then be used to provide chapterised versions of the linear output, make individual stories findable by search and available as items on web pages. Ultimately, by using Object Based Media (OBM) techniques these stories can be recombined to provide personalised and/or interactive experiences.

This paper begins by outlining our hypothesis that audiences will gain more value from our linear output by giving access to individual stories and items compared to entire, undifferentiated programmes and outline the user testing that supports it. It explains the approach being taken to recover programme metadata and how that can help achieve content segmentation at scale.

The paper then describes how OBM techniques are being used to automatically combine segmented content to make new interactive and personalised experiences, including the prototype code base and the workflows developed so far. The paper concludes with a summary of the key findings from the work to date, including audience benefits, data consistency and the role of editorial oversight.

### **INTRODUCTION**

Radio and TV news and magazine programmes are built up from a series of components brought together to create a singular linear programme which is broadcast and available through on demand services. News and magazine programmes include many different stories in a single episode. However, if there is no way of highlighting these items and making them discoverable following the linear broadcast, then the value of these individual stories can remain locked within these programmes and hidden from the wider audience.

By recovering these stories as individual objects through segmentation, we can create the opportunity to build new forms of personalised and interactive experiences. However, to build these on a cost-effective basis there is a need to curate and schedule content items with minimal human intervention. Ultimately, by creating editorial templates for personalised versions, automation can then be used to find the best fit of content to the individual's needs based on metadata that describes the items and the experience formats, comparable to music scheduling in most radio stations.



## THE MODULAR CONTENT HYPOTHESIS

Our hypothesis is that we can extract far greater value from the BBC’s world-class linear content for our digital platforms if we can provide access to it at the level of the individual story. Money spent on the linear output does not currently translate to a rich, accessible digital experience and sells itself poorly on the BBC’s digital platforms. This amounts to valuable content and world-class journalism being locked away from the audience. BBC R&D’s News Labs has been exploring ways in which we can improve the digital currency of the BBC’s linear content. Currently BBC flagship news programmes are offered as entire programmes on BBC Sounds and iPlayer with metadata that describes the programme as a whole. For example, a typical "billing" synopsis for Radio 4’s *Today Programme*, the BBC’s main speech radio breakfast show is billed as “180 mins of News weather and sports desk” with no further detail. More information should foster better search engine optimisation and audience engagement, and make the audience feel closer to the content, because there is richer detail of all the stories covered in the programme.

### Production metadata

In contrast to the sparse audience facing digital offerings, the production cycle producing linear content includes thousands of hours of effort that goes into building detailed and descriptive data for each programme in the form of their running orders. These running orders are used to define, structure and control the production of all of BBC’s linear News (and some Sports) output. They contain rich descriptions of the items in the transmitted programme and extend down to the level of the individual story and scripts. However, there is currently no mechanism for this data, or at least an audience-facing version of it, to reach the BBC’s audience and make the individual stories searchable within a programme.

It is also worth noting that problems exist with concepts such as *title* and *script* or *synopsis* as they have different requirements in a production context from an audience-facing publication context. An issue we will return to later.

### The Research Question

News Labs set out to answer the question, can existing data and processes support:

1. Digital artefacts that deliver greater value across more platforms?
2. Curated experiences tailored for each member of the audience?

The work has been framed in four steps towards personalised content, Figure 1.

1. Gathering existing 'story based' production data, in order to...
2. Provide versions of programmes that are chapterised by 'story' but also supporting...
3. Editorially coherent, free-standing stories ('best bits') combining digital native content with items from linear output and perhaps surfaces as recommendations, and...
4. Semi-automated or fully automated content curations which is personalised and enables user choice across the whole breadth of BBC linear and digital output.

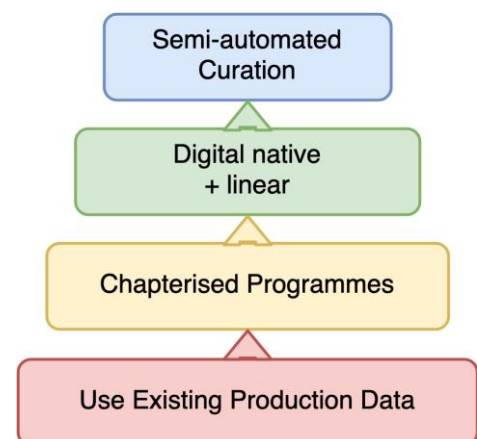


Figure 1 – The steps towards personalised content creation



It is this fourth and final stage where the Object Based Media team from BBC R&D are working, turning individual items into personalised and interactive experiences.

### User Research – Chapterisation

User research was conducted using the BBC World Service - *Global News Podcast* (1) to evaluate the benefits of chapterisation for the audience. This is the BBC’s most popular podcast, published twice a day on weekdays and daily at weekends. News Labs were approached by the production team who wanted to add chapters to the podcast. A test was run to create an episode of the podcast using production data to drive chapterisation and provide a segmented version of the podcast as a web page.

The research was conducted with around 2000 participants from the World Service Global Minds survey group (2). Of these, 89% found the functionality useful, 62% listened to the test podcast in a different order to the linear offering and 72% of respondents said they were more likely to consume a podcast with this functionality. This validates the hypothesis that adding chapters to linear content can provide a significant benefit for the audience as well as increasing its reach.

### SEGMENTATION AT SCALE

#### 1. News Labs Slicer

The question then arises as to whether it is possible to provide segmentation of linear programmes at scale? News Labs built a prototype called Slicer AV to do exactly that, Figure 2. Slicer AV harvests and stores programmes running orders as they are output from the BBC's Open Media Newsroom computer system. The running orders contain scripts that the presenters use to introduce each story within a programme. The data from Open Media is standards based so our approach is applicable to any new or sports production that uses a Media Object Server (MOS) based running order system. The MOS Running Order Manager has been made available as open source (3).

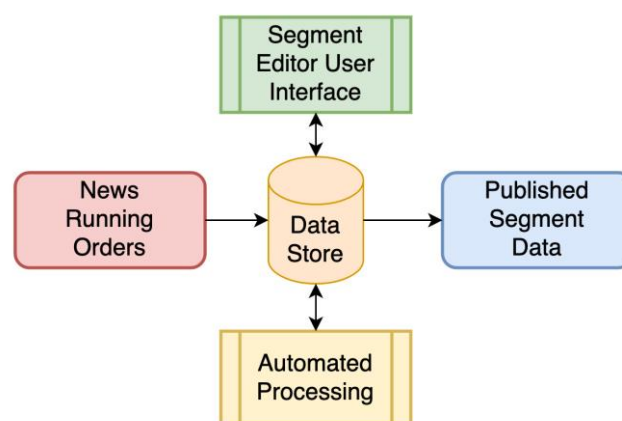


Figure 2 – Segmentation workflow

Slicer AV also transcribes the programme's audio using BBC R&D's speech to text engine based around Kaldi. It then searches the programmes transcript for the occurrence of the phrases within its running order scripts (or something like them). In this way it is able to determine when these words were spoken within the programmes timeline and mark out story chapter start points. The success of this alignment processes naturally depends on the quality of the transcript and the extent to which the presenter sticks to the scripts in the running order. Some programmes auto-segment really successfully while others, in particular those that are more free form in style, are more challenging. The auto-alignment process is an area of ongoing research within the News Labs team, and we hope to explore and test different models soon.

A web-based segment editor has been created to allow a producer to check and tidy up the segment timings, edit the segment titles and provide a suitable, audience facing synopsis as well as adding tags from the BBC Things classification system (4) and marking items for publication where appropriate. Not all items in a linear programme are



re-usable outside the context of the podcast or broadcast. For example, the opening and closing items and any item referring to other parts of the programme will make little sense as stand-alone clips, so whilst they may be valid for chapterisation, these items are not suitable for reuse.

### **Chapterisation prototype for BBC Sounds.**

News Labs has created a plugin for the BBC's Standard Media Player (SMP) - the media player used in BBC Sounds and BBC iPlayer pages - to surface chapter-based navigation in BBC Sounds Web pages. The plugin went live in April 2021 for a 6-week trial on the Sounds Web pages for the BBC Radio 4 magazine programmes *Woman's Hour* and *You and Yours* and for the World Service *Global News Podcast*. The plugin features a number of new metrics on per-chapter clicks and can map those back to the programme running orders to track the interest in each story.

News Labs supported the Sounds metadata team to slice up episodes for these programmes and publish them on a daily basis. The scope of the trial is limited to the Sounds Web version, which means the chapters do not work in the Sounds mobile app or TV app. The data collected (and user research being conducted by Sounds during the latter half of the trial), is being used to help decide if chapter-based navigation should be made more generally available, and develop versions for the Sounds apps on iOS, Android and TV devices. This is the beginning of a process of giving the audience deeper access to, and a better understanding of, BBC content. It may also lead to better recommendations, at the individual story level as well as the programme episodes.

## **2. Internet Fit Radio**

Another BBC initiative to improve metadata for content reuse named Fenchurch was developed by the Internet Fit Radio team, now part of News Labs, Roberts and Eldridge (5). The system runs across the BBC's popular music and speech radio stations; BBC Radio 1, Radio 2, 1 Xtra and 6 Music. It captures the presenter's desk and playout system events along with associated metadata. It uses speech-to-text on the speech segments and the text is fed into BBC R&D's automated tagging system called Starfruit which adds tags according to the BBC Things entity system (6). Further integration with other internal data systems enables the identification of the individual music items. The data from Internet Fit Radio is somewhat different from the Slicer output as it very accurately collects timed events and music track identities but does not contain script data. Also, because these are presenter-lead music programmes the speech overlaps the music content both at the beginning and ends of music items and presenters will talk over music beds. This makes the speech to text conversion less accurate and poses problems in trying to identify clean speech items for re-use.

## **3. Wildlife AI Camera Feeds**

BBC Research & Development has also been investigating how artificial intelligence (AI) technologies such as computer vision and machine learning can be applied to assist in the production of natural history programming. BBC R&D has been collaborating with the BBC's Natural History Unit (NHU) to test systems developed to help monitor some of the remote wildlife cameras for the BBC wildlife programmes *Springwatch*, *Autumnwatch* and *Winterwatch*. The system monitors the cameras for activity and logs when something happens, attempts to identify the animal and makes a recording, Dawes (7).



The first part is an object recognition system that is designed to detect and locate animals in the scene. It is trained to recognise both birds and mammals and tracks how long an animal is visible in the scene. The images are passed on to a second machine learning system which attempts to classify the animal's species using a convolutional neural network. If the system is sufficiently confident of an animal's species it will label the photograph and also add the label to the user interface and the confidence score.

The system was developed as an aid to programme production teams to enable them to more easily find interesting activity on a particular camera's video. However, these video clips also provide a rich feed of interesting content for repurposing in other ways and have the advantage over news stories of being timeless, so more suitable for ongoing experimentation with different formats.

## **CREATING INTERACTIVE AND PERSONALISED CONTENT**

This section explores how we view the task of taking existing linear content and reformatting it into personalised versions or interactive experiences and the following sections will focus on the practical steps to achieve those processes.

### **The role of templates in linear content**

News and magazine programmes on radio and television are usually built to a regular repeating template, encapsulated in the running order, which shapes the ebb and flow of the programmes from the high importance headline story at the beginning through to the lighter, more uplifting “and finally” story to end, often followed by a weather forecast. There are also items such as idents that provide signposting. Whilst individual stories change with each programme, the overall structure is usually stable and only changes for major events. In the case of live programmes, particularly with radio, this structure will run to a fairly strict set of clock times. Indeed, in the case of radio breakfast programmes the audience often uses the programme in place of a clock and will complain if regular items get moved around. To some extent this structure is useful in helping to deconstruct the programme into the component parts for chapterisation and reuse whilst also pointing to a way of using a regular structure to build a personalised and/or interactive versions.

One of the key things to understand when re-using segments of a programme is their function and context and how that function might be altered when recompiling a new version. The opening of a programme will usually contain the name of the programme and the network, often with a time check and a summary of the stories to follow. If we are creating a digest from a programme series, the identity of the programme is still relevant, the network less so, whilst the contents of the new version will be different, and the time check will be completely misleading. Similar issues exist with other items of signposting like trailers and idents. However, there is still a need to give the new versions a sense of pace and timing, breaking up stories with equivalent items to give some thinking time before the next item starts and signposting progress through the new version. The structure within a new linear version, such as a podcast can reuse broadcast techniques.

### **Formats for interactive content**

For interactive content, structure and signposting are a greater challenge, as indicated by the audience feedback to previous OBM experiences, *Instagramification* and *Click 1000*, Armstrong et al (8), along with discussion of the user experience of the *Radiophonic Travel Agency* experience, Brooks (9). With a linear personalised version of a programme the usual navigation affordances are available, particularly the timeline indicating progress





through the media. Further signposting adds to that without the audience having to refer to the timeline. By contrast, if the length of the content varies according to the choices made by the audience, then a simple timeline is not sufficient. This is a largely unexplored topic where it is difficult to extrapolate from a small number of items. However, it is clear that some form of wayfinding and progress indication, perhaps similar to maps and achievements in video games, would significantly improve the user experience and reduce the fear of missing out on interesting items, Glancy et al (10). As more of this style of content is produced, we will gain a better understanding of what templates are likely to be successful for interactive content and by automating processes in the production of such content we can explore these ideas more rapidly.

### **Object-Based Media Tools and Data Model**

We are building on our StoryKit suite of software tools for authoring and delivering object-based programmes. We have added tools for analysing the experience structure and visualising audience behaviours (8,11) and have made the tools available to a wider range of creators through BBC R&D's Makerbox portal (12). At the core of these tools is a common data model describing the media items, how they are displayed and how they interact with the audience. The data model and StoryPlayer from BBC R&D's StoryKit toolkit are now available as open source (13,14).

The main user interface in StoryKit is the web tool StoryFormer which provides general purpose, graph-based authoring. It enables teams to build stories that are object-based from the outset, making adaptive narratives and interactive experiences as easy to produce as linear programmes. It also has the potential to enable new functionality for linear programmes, allowing production teams to describe their content at a high level of granularity. StoryKit was used to create *Instagramification*, *Click 1000* and *Discover Your Dæmon*. However, all three productions were still quite labour-intensive processes and unsuitable for repeated regular content production such as a daily podcast.

We then began a new stream of work looking at how new experiences could be created more easily by constraining the production process to a template or even automate it altogether. The feeds from News Labs Slicer, Internet Fit Radio and the AI wildlife cameras provided the content feeds with metadata on which we could build new content.

### **EXPERIMENTS IN AUTOMATION**

Our automation workflow was developed through initial experiments to find the types of experiences which could be built by automation. The most obvious starting place was to borrow the ideas from the News Labs "Skippy" prototype which became part of the BBC Voice service, Lei He (15).

#### **Skippy**

Skippy provides the ability to jump from the current story to the next story or to the next section or programme. The initial test took a number of episodes of the *Global News Podcast* and the data feed from Slicer to create an OBM version where there was an opening menu allowing the episode to be selected. The items from the programme were presented in the same order as the original podcasts, but with the ability to skip to the next item, skip to the next programme or return to the menu. This turned out to be a useful tool for checking the success (or otherwise) of the segmentation and introducing basic concepts but did not make a particularly convincing demonstration or user experience as it did not significantly improve on simple chapterisation.



### Top tags / story catchup / bird selector

The second iteration used tags to group stories by topic and select the most frequent topics to present in the experience, with each group of stories presented in temporal order and included the ability to skip forwards. The initial work on *Global News Podcast* presented the top topics from the past week. This provided a more compelling demonstration of the potential of this work, often picking up a story and showing its development, such as the USA presidential election, or contrasting stories from different parts of the world on individual topics, particularly the impact of Covid19.

Discussions with a journalist from News Labs highlighted the problems of loss of context, particularly the date of the story, where the understanding of events has changed. Because the metadata included this information the demo was easily modified to provide text overlays adding the programme title and date of publication. This also enhanced the feeling of following a developing story as you could now see the story developing in time. The topic and item number were added to improve the navigation along with the tags for the story, Figure 3.



Figure 3 – Prototype interactive podcast.

This approach also produced usable demonstrations with other programmes, particularly Radio 4's *You and Yours*, where there is a good clustering of tagged topics and reasonably accurate segmentation. This code also proved to be usable across the different media and data sources. The same process applied to the Internet Fit Radio data proved successful in grouping together content, where clusters of tags occurred. However, the more challenging nature of the content, with speech overlapping music and issues of rights with music tracks make these demonstrators less viable.

This was also the case for content from the *Autumnwatch* bird feeder cameras. Here, by selecting clips tagged as containing single bird identifications at high levels of certainty the code produces a bird watching experience where you can select which species to view. This simple, but effective demonstration lead onto the idea of using the same set of clips to create a bird identification quiz.

### Automated Quiz Generation

The quiz was developed by selecting a total of thirty-two video clips, eight clips for each of four bird species. For each question four clips were used, one for each of the species, with one acting as the question. As the question video is displayed the user has to select the name of the species from a choice of four. If the user selects the correct answer the video is repeated saying the answer was correct. If the user selects a wrong answer they are shown a video of the species they selected with the overlay saying the answer was incorrect. A running score is displayed and the overall result given at the end. A manually enhanced version, using footage from the 2020 *Autumnwatch* camera feeds was published by BBC R&D in advance of the 2021 *Springwatch* series (16) and a second quiz was manually built and published as part of the *Springwatch* offering, using feeds from the nest cameras (17).



### **Templated Quiz Tool**

The examples above show how scripts can be applied to well-marked-up content in order to create personalised or interactive experiences. In these cases, human intervention may be necessary at the segmentation or tagging stages to ensure that these have been achieved with sufficient accuracy and precision; the transformation of these segments into audience experiences, however, is fully automated. Another approach is to surface the segments to production teams in a lightweight, but highly specific tool, in which they are assembled into experiences. Our initial exploration in this area was to generate a quiz-building tool. Users specified questions and answers in a simple Web form, associating media with each; the quiz was held in an intermediate format that had a template applied to transform it into an instance of the OBM data model (this two-stage process allows for different flavours of quiz to be generated from the same interface; applying one template might create a fast-paced quiz while another might be more suited to learning). This manual approach to assembling experiences is a half-way house between the near fully automated systems described above and handcrafting in StoryFormer. It is much simpler and faster than the latter as the interface is highly tailored but is more tolerant of errors in segmentation and tagging than the former. This type of tool is not restricted to building experiences from scratch, however; it was also designed for reviewing and editing automatically generated content and since the output of the tool is in our standard data model, the quizzes can be viewed and edited in our StoryFormer tool, allowing them to be customised to a fine degree or to be incorporated into a larger experience.

### **AUTOMATION CODE BASE**

In order to automatically create usable experiences from segmented content and the raw descriptive data, the workflow was broken down into a number of separate processing steps, Figure 5. Note all the media content is streamed from its Uniform Resource Name (URN) or Uniform Resource Locator (URL) so no media processing takes place, only the manipulation of metadata (e.g., setting in and out times to trim a segment)

### **Pre-processing and tagging**

Each of the three sources of data are in different formats and different levels of completeness. Where slices have not already been manually tagged, the script adds tags to slices by passing the script data in the running order through Starfruit to create machine derived tags. The case of Internet Fit Radio, the data already contains tags generated by Starfruit for the speech content, but additional tags are added to identify the item type - speech, music, insert, trail, news and the timings are resolved between different versions of the music data. Springwatch data is reformatted to make it conform to the Slicer intake format and tags are created by using the bird identification data for single sightings. Multiple sightings are currently discarded.

### **Common format creation**

The following stage conforms the data to a common list format with consistent nomenclature. Each contains a list of items relating to programme episodes with each programme represented as a list of self-contained items. Each item now contains the programme title and media URL or URN so that they can be used separately by downstream processing. At this point, all data sources have been conformed to flat lists of tagged time intervals with domain specific tags.



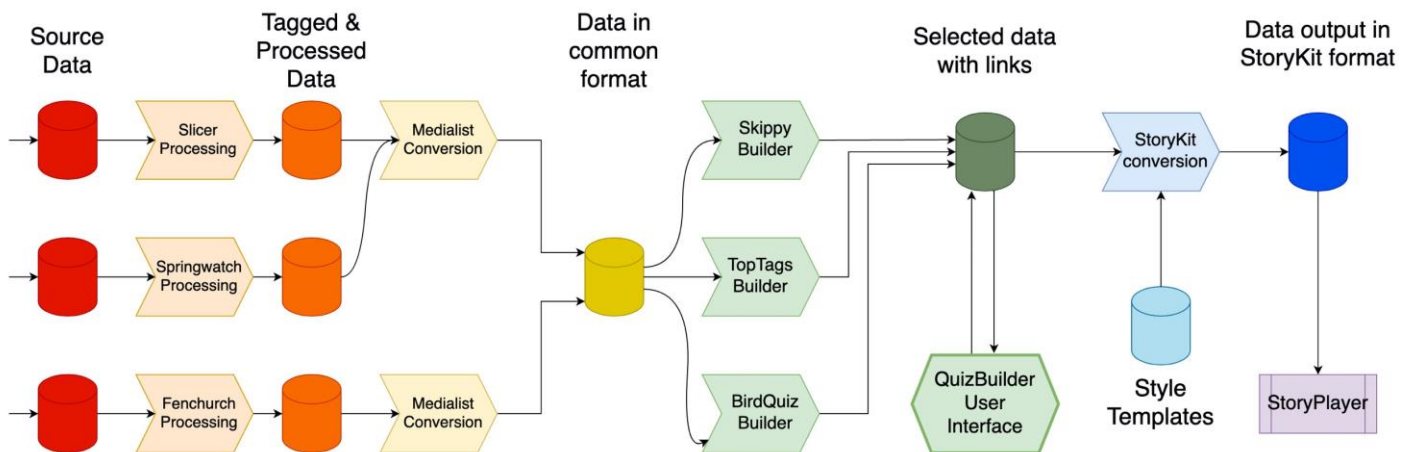


Figure 5 – Segment metadata processing workflow

### Item selection and linking - format creation

This is where the items are selected for downstream use and the items are arranged with the addition of the inter-element links. At this point, *introduction*, *menu* and *interstitial* items are added. Items are also labelled according to which template should be used to add behaviours at the next stage. This is a generalised representation which borrows the link representation from the OBM schema and adds layout information for StoryFormer but is otherwise a list of media items and descriptive metadata.

The simplest script creates the “Skippy” version. It takes any media list and creates grouping of data for items from each programme list such that you can jump to the next item, to the next programme or back to the menu. It does some tidying up to remove overlaps between items.

The more complex script creates the “Top Tags” experience. It counts the number of times each tag is used and sorts the data into groups. It then removes duplicates, selects the most used tags and links these items under each tag in chronological order. It also adds introduction, menu and end items. This is a general-purpose script that works on any data with a reasonable number of re-used tags.

Finally, the “Bird Quiz” is a more specialised script is currently specific to the wildlife camera data. It looks for a set of birds that have each been identified in 5 or more items and builds a basic bird identification quiz from them. One clip is selected from each set to be the question and the others are then the wrong answer, if the answer is correct the clip is repeated. It adds an intro and interstitials and an end item. The Quiz Builder web UI tool can import and export the version of this data in the form of a quiz and create new quizzes.

### Conversion to OBM Schema representation

The final script takes the linked items and conforms them into the OBM Schema. It uses style templates to add behaviours to each item based on their grouping, such as *correct* and *incorrect* answers in the quiz layout. This provides a compact way of describing how an item should be displayed and the functionality for scoring a quiz and ensures consistency between items. The resulting files can then be imported directly into StoryFormer and StoryPlayer.



## KEY FINDINGS - SO FAR...

This work on segmentation and content reuse is still at a relatively early stage, but a number of issues have surfaced from this work.

### Audience benefits

The user research on chapterisation clearly demonstrates that the audience not only like having chapters available in a programme, they actively made use of it as the majority of people surveyed listened to the items out of order. This indicates how additional data and improved navigation of linear on-demand content can significantly improve the user experience. Further benefits from making these stories discoverable and reusable within personalised and interactive offerings will build on these. Whilst further trials and user research is being carried out to quantify the impact, it is clear that this approach is releasing additional value from linear content.

### The role of editorial oversight

Even with good quality segmentation and metadata, it is clear that there is still a need for editorial oversight before the publication of each item. Issues arise with production data being unsuitable for public consumption, segmentation not finding the most appropriate *in* and *out* points and content items which cannot be used outside the programme context. A producer checking each item can verify that items are viable only within a chapterised version of the programme or as stand-alone items that can be reused in other contexts and can ensure that permission has been given for re-use. They can also deal with issues where the original item has been withdrawn and/or replaced with a re-edited version.

### Production data vs publication data

One of the problems with production data is that the title of an item in a running order will often be unsuitable as an audience-facing title. So, a story may simply have the title "Russia" in the programme running order, but this is not sufficient for publication. The title may also contain extraneous production information, such as the item's length or the journalist's name, for example "*RUSSIA-Reporter Name R4 18*" for a story which came from the Radio 4 1800 news programme. Similar problems arise with script information containing instructions to the presenter and information about interviewees. So, scripts need to be reviewed and edited if they are to be reused or rewritten as a short synopsis.

### Production data consistency

Because the segmentation process relies on finding a match between running orders, script information and the audio-visual content, problems arise where production processes are inconsistent. Changes to item names, instructions in scripts and items being out of order all disrupt the segmentation process. While linear production remains separate from the digital offering, production practices will remain focused on the linear output.

### Tagging issues

Automated tagging systems for news and sport rely on machine learning being used on a regular basis to update the model with new tags that are being created in response to new events. This means that any automated tagging system will lag behind the tags being employed by journalists and the most important breaking stories will be the least well served because of insufficient training data. This gap can only be filled by human oversight, associating the item to new tags as they are being added to the system.



### **Responding to major events**

Recent events in the UK have demonstrated how entire broadcast schedules can be overturned by major events. Such major news stories not only disrupt the broadcast schedules they also create sensitivities around certain topics. Human oversight of publication is essential in times of major news stories, quite apart from the failure of any segmentation system because of last-minute running order or schedule changes.

### **Other data related issues**

This work has confirmed the need for non-destructive editing as far as possible. As one process flows from another, if data is over-written then it is difficult to debug when the evidence has been destroyed. Also, the original data may prove useful for downstream processes. Other challenges have arisen when trying to reconcile different data sources describing the same content but in different contexts. A set of data which is used to create items on the programme web page may have additional metadata about the contents, but lack correct timing information, so a multistage process is needed to reconcile this with the timed data. Finally, there is the issue of compatible data sources. In order to find additional content to enhance a story we need to be able to access information across a range of sources which are likely to use different and sometimes conflicting ontologies.

### **Personalisation**

This work has yet to address the challenge of individual personalisation of these experiences. There are significant challenges here around use of each individual's personal data which need to be addressed in-line with the BBC's public service values. This is the topic of another ongoing research theme at BBC R&D (18).

### **CONCLUSION**

This work is the first step towards unlocking the value in the BBC's linear content and in unused footage, making it available to a wider digital audience. This work requires good data stewardship and artificial intelligence with machine learning techniques to be partnered with editorial oversight to create stories and clips that are suitable for publication. The work has developed to the point where public facing trials are underway with chapterisation for three BBC programmes and further work is ongoing to improve the reliability of the segmentation as a better understanding of the way production scripts are used is developed.

Work on creating interactive and personalised content from this segmented material is still at an early stage. A modular code base has been developed which can automate the creation of demonstration experiences and new versions can be developed and refined with relative ease. However, more work is needed to understand the editorial challenges around news and magazine content and the types of experience that will bring value to new audiences. Early public facing demonstrations have been published using the wildlife camera feeds building on the ongoing collaboration with the programme.

The techniques outlined are beginning to demonstrate both the opportunities and challenges in re-using linear content feeds and data and the benefits that can bring to both audiences and broadcasters.



## ACKNOWLEDGEMENTS

We would like to acknowledge the contributions from Mark Langton and Ben Nuttall from News Labs and Matthew Brooks, Andy Jones and Craig Wright from the OBM Toolkit team along with those who have contributed to these projects over many years.

## REFERENCES

- (1) Global News Podcast. <https://www.bbc.co.uk/programmes/p02nq0gn>
- (2) World Service Global Minds survey group. <https://bbcglobalminds.com/>
- (3) MOS Running Order Manager. <https://www.bbc.co.uk/opensource/projects/mosromgr>
- (4) BBC Things. <https://www.bbc.co.uk/things/about>
- (5) Chris Roberts and Luke Eldridge, June 2020. An intelligent approach to capturing and using radio metadata, EBU tech-I, Issue 44. <https://tech.ebu.ch/docs/tech-i/tech-i-044.pdf>
- (6) Jeremy Tarling, June 2020. How the BBC is approaching metadata quality measurement, slides for ISKO UK Meetup. <https://www.iskouk.org/event-3863672>
- (7) Robert Dawes, June 2020. Using AI to Monitor Wildlife Cameras at Springwatch, BBC R&D Blog Post. <https://www.bbc.co.uk/rd/blog/2020-06-springwatch-artificial-intelligence-remote-camera>
- (8) Michael Armstrong, Sally Bowman, Matthew Brooks, Andy Brown, Juliette Carter, Andy Jones, Max Leonard, Thomas Preece, June 2020. Taking Object-Based Media from The Research Environment into Mainstream Production, SMPTE Motion Imaging Journal. <https://ieeexplore.ieee.org/document/9109608>
- (9) Matthew Brooks, December 2020. Take a break with the Radiophonic Travel Agency, BBC R&D Blog Post. <https://www.bbc.co.uk/rd/blog/2020-12-radiophonic-workshop-interactive-sound-story-travel>
- (10) Maxine Glancy, Lauren Ward, Nick Hanson, Andy Brown, Michael Armstrong, Sep 2020. Object-Based Media: An Overview of the User Experience, White Paper WHP 390. <https://www.bbc.co.uk/rd/publications/whitepaper390>
- (11) Lauren Ward, Maxine Glancy, Sally Bowman, Michael Armstrong, September 2020. The Impact of New Forms of Media on Production Tools and Practices, IBC2020. <https://www.ibc.org/download?ac=14535>
- (12) <https://www.bbc.co.uk/makerbox>
- (13) OBM data model. <https://github.com/bbc/object-based-media-schema>
- (14) StoryPlayer. <https://github.com/bbc/storyplayer>
- (15) Lei He, April 2020. How Skippy made it from a hack to your home, BBC News Labs. <https://bbcnewslabs.co.uk/news/2020/skippy-hack-to-home/>
- (16) <https://www.bbc.co.uk/rd/blog/2021-05-springwatch-bird-quiz-machine-learning>
- (17) <https://www.bbc.co.uk/programmes/articles/14nPS3zsrdRsbt86SSfp6Wm/whose-chicks-are-these>
- (18) Personal Data Stewardship: A new way to store and use personal data. BBC R&D. <https://www.bbc.co.uk/rd/projects/bbc-box-personal-data-management-privacy>