SHREWSBURY MUSEUM & ART GALLERY AT THE MUSIC HALL

Insight 2

Intelligent Exhibit
Proof of Concept

A collaboration between Black Radley, Microsoft and Shrewsbury Museum

12 April 2017

Black Radley

Joe Collins **Black Radley Ltd**
joe_collins@blackradley.com

www.blackradley.com
T: 0845-226-0363

Table of Contents

Contents

Executive Summary

1. Museums and galleries already make use of audio visual technology. This technology can be enhanced so that artefacts and exhibits can respond intelligently to visitors, providing a customised narration.
2. A “proof of concept” prototype was produced during a one week collaboration between Black Radley, Microsoft and Shrewsbury Museum.
3. The availability of cheap and fast facial recognition software allows this to be done for under £100 using free (Open Source) software.
4. Updating the narration could be relative quick so it can be placed in the control of local subject matter experts who can change the narration when required or in response to visitor feedback.
5. Such a device can also gather information about its own performance. So that different narrations and exhibits can be evaluated and the behaviour of visitors can be understood.

1 Visiting Museums and Galleries

The nation's museums and galleries contain a rich store of original and authentic artefacts. Visitors seek to experience and understand those artefacts. Understanding the back story to an artefact can make it "come alive" and acquire meaning to the visitor.

Ideally each visitor would be able to experience the gallery or museum with its curator. Curators are the experience experts, able to reveal the meaning and context for their collections. They can adapt the narrative to the needs of the visitor, perhaps not saying anything when the visitor has had enough. Unfortunately expert curators are a limited resource so museums and galleries have turned to audio visual technology to help out.

1.1 Audio Visual Technology in Museums and Galleries

Museums and galleries are continually looking for ways to bring their collections alive. Using a variety of methods and channels so that visitors can experience the artefact according to their own personal style. The painting below from Shrewsbury Museum and Art Gallery is an excellent example.

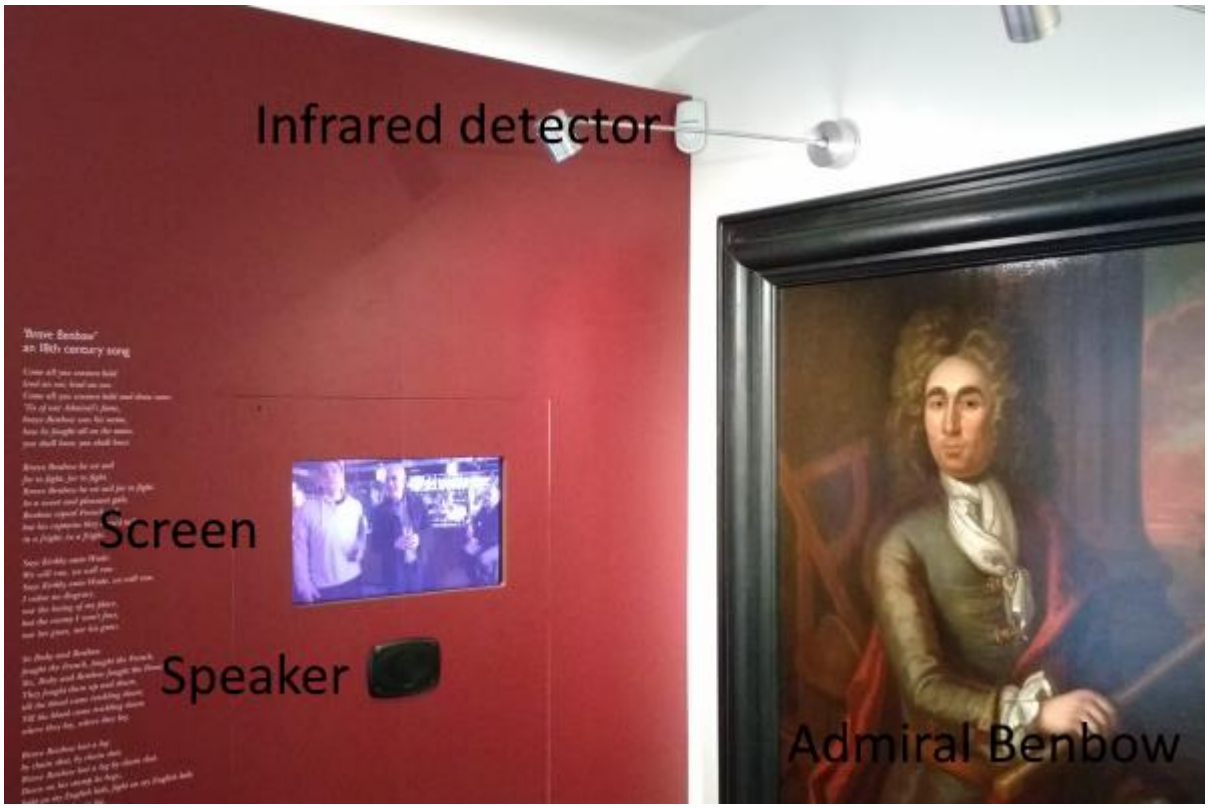

By the side of the painting of Admiral John Benbow (a Shropshire hero) is an audio and visual presentation together with text. The audio and visual presentation is

triggered by an infra-red detector. Having been triggered the audio and visual presentation runs for about 5 mins so there are seats for visitors to sit and watch. The presentation continues even if the visitor turns away or goes to visit other exhibits.

1.2 An Adaptive Experience

Ideally each visitor would have their own curator guiding them around the museum or gallery. It is not yet possible to replicate this experience using technology but it is possible to go some way to providing a customised and adaptive experience with the exhibit reacting to the visitor.

Image and facial recognition have been available for some years now. It is possible to use a camera identify visitors and deliver a narrative adapted to the visitor and where they are looking. Some systems have been produced but until recently this has been prohibitively expensive and required specialist equipment. Now the cost of such technology has dropped substantially, sufficient to bring it into the range of hobbyists, makers and DIY enthusiasts. It should now be possible to create a device that can recognise and respond so visitors using free (Open Source) software and hardware costing less the £100.

Such a system could recognise when a visitor is looking at an artefact or exhibit and be able to provide a different narrative based on the visitors age. It should also be able to recognise when the visitor is no longer looking at the exhibit, even if they are still in front of it.

2 Proof of Concept

A combined team from Black Radley Limited, Microsoft Corporation and Shropshire County Council came together at Shrewsbury Museum and Art Gallery, for one week in April 2017, to confirm this proposition and to produce a working “proof of concept” prototype.

Shrewsbury Museum has a wealth of artefacts that could be used for the “proof of concept” but it was decided to focus on the earliest known panorama of Shrewsbury, from 1630. The painting shows many of the buildings that are still present in Shrewsbury today, but from an unusual and implausible perspective. The history of the painting itself is interesting, until recently it was hanging on the kitchen wall of a farm house in Wales before it was donated to Shrewsbury museum.

The mounting and positioning of the painting allows a device to be conveniently placed under it. None of the equipment used was very bulky so it would be entirely possible to place a camera inside a display cabinet. However, accessing cabinets would have inconvenienced museum staff so an artefact with a convenient surface beneath it was selected for the prototype.

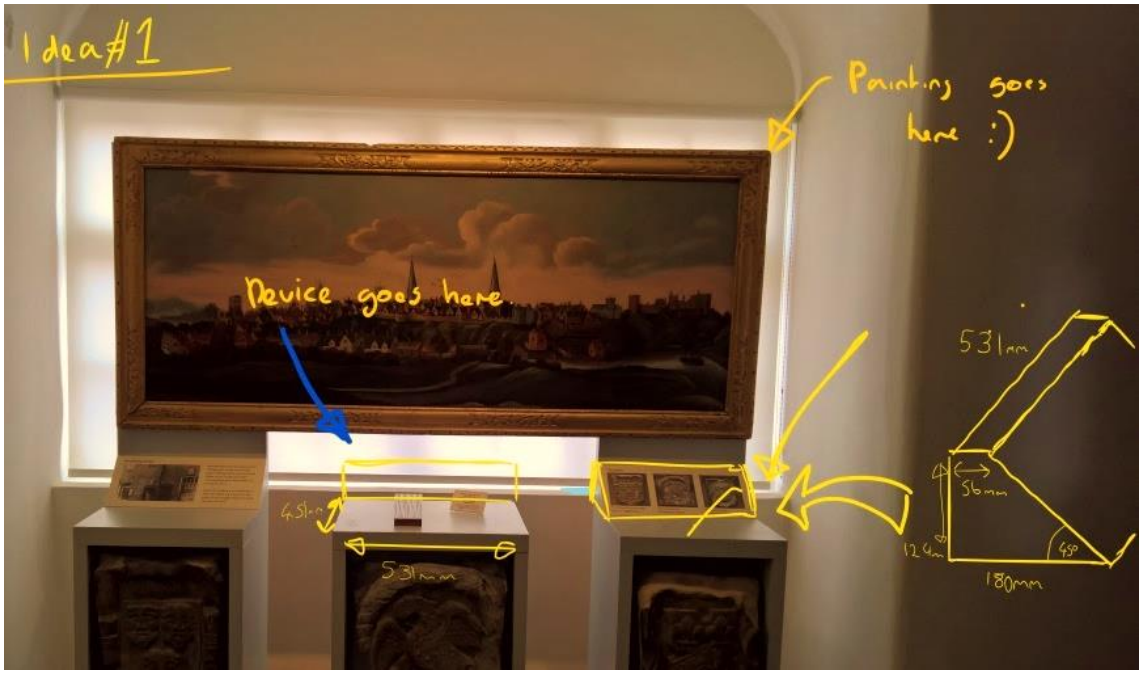

The prototype used conventional pieces of hardware. In fact, no hardware was purchased for the project. All of the components were borrowed or scavenged from the cupboards and desktops of the members of the team. This was a great advantage since there was no budget for the project and everyone was donating their time for free.

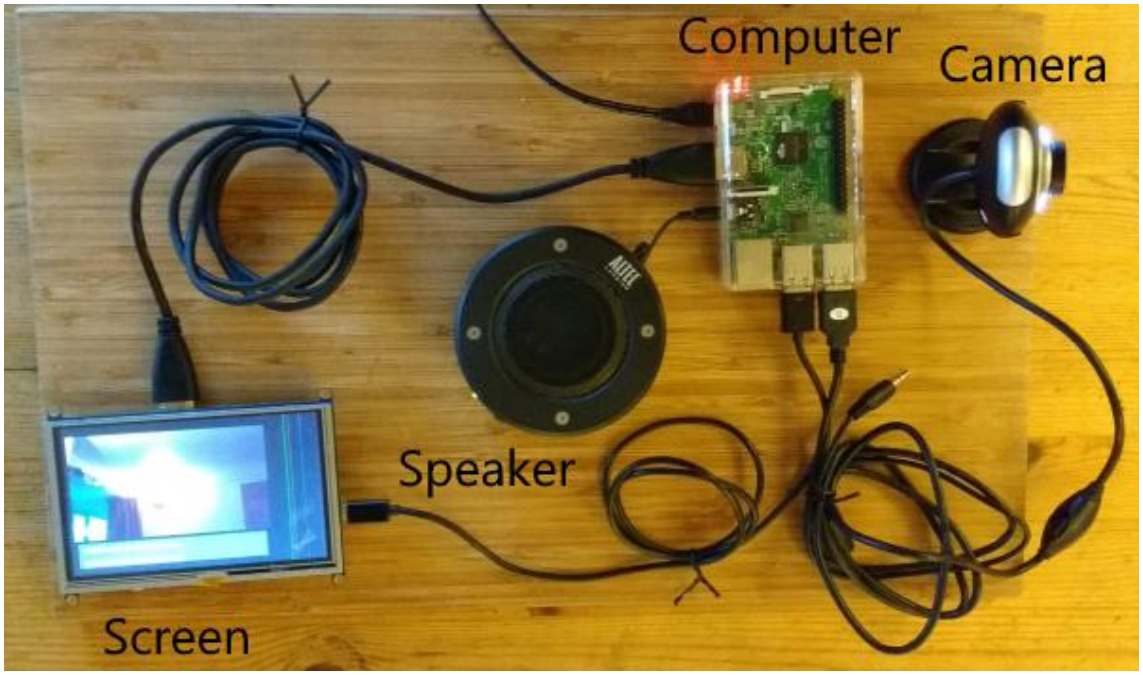

The computer used was a cheap computer of the sort used by schools to teach computer programming and information technology (£50). The camera is a conventional low resolution webcam, used for video conferencing (£20). The speaker is of the type used to plug into an MP3 player (£20). The screen was not really necessary but was useful while the team was working on the device.

The hardware was concealed in a box made out of cardboard. It is not perhaps the most beautiful container, but the dimensions are being used to construct a more professional box that can be blended in with the current surroundings.

All the software used in the project Open Source and freely available at https://github.com/blackradley/dinmore.

2.1 Facial Recognition and Narrative

The device was able to distinguish between the arrival of groups or individuals. It was also able to categorize each visitor by approximate age. Groups were greeted with "Hello Everyone" and individuals were greeted with "Hello there". The narration was split into separate sentences so if the visitor turned away from the painting or left then the narration would stop playing at the end of the sentence.

If anyone viewing the painting was below the age of 20 years then a description of the painting was provided orientated towards younger visitors. If everyone was above 20 years old a different narration was provided. Whilst determining age and delivering a different narration is clearly achievable, it does raise the question as to whether this is effective or desirable. Fortunately the device records information about visitors so that any response to this question can be based on information about how visitors behave and how they linger in front of an exhibit.

2.2 Performance Information

As a by-product of the facial recognition process the device also recorded its estimates of age, gender and facial expression. It is therefore possible to determine how long a visitor spent viewing the painting. Using this data it would be possible to determine if the narration was affecting the visitor's dwell time and presumably enjoyment of the painting.

The data also indicates which visitors are returning visitors by attempting to match previously seen faces with new faces. This process is not as exact as vehicle number plate matching so the device also records a level of confidence as a percentage. For example wearing glasses is assumed to be a fixed feature of a face, so taking off your glasses means you are assumed to be a different visitor.

The information is available in a conventional spreadsheet for study by curators and managers.

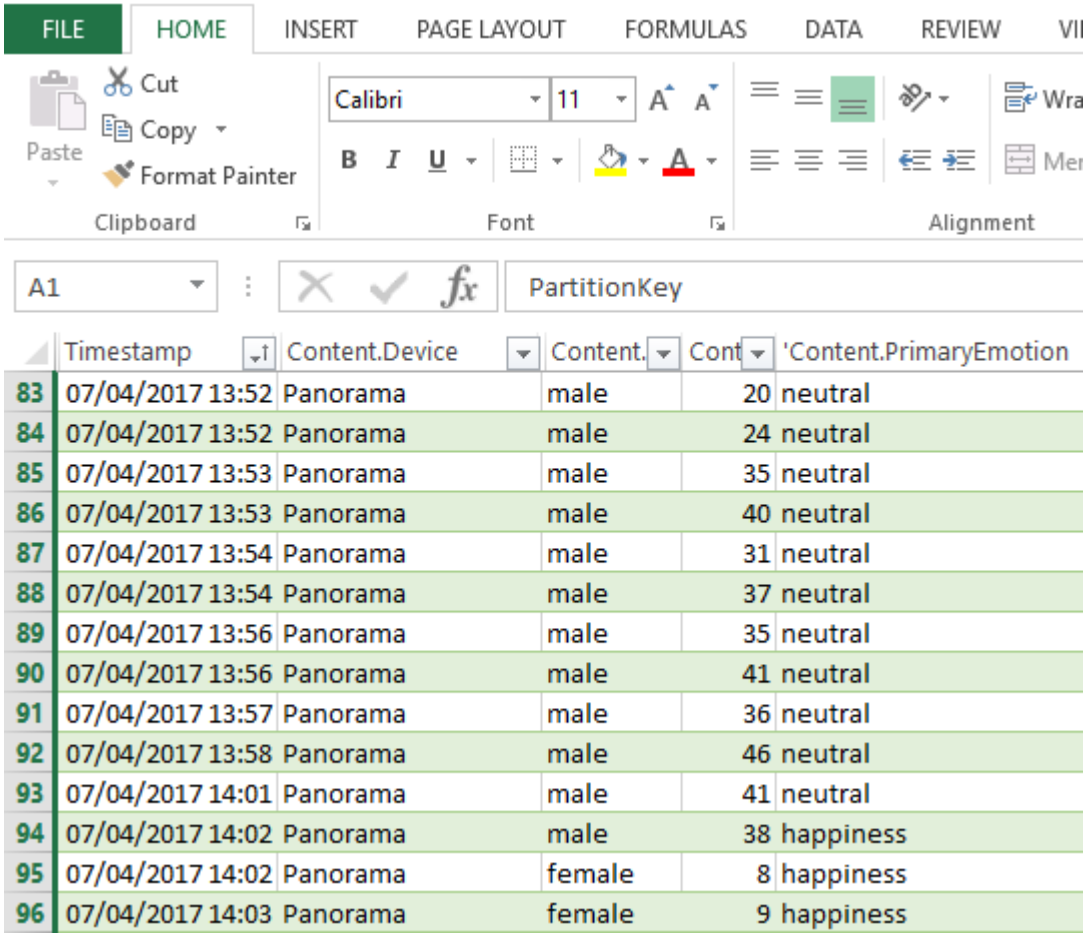

	Timestamp	Content.Device	Content.	Cont	'Content.PrimaryEmotion
83	07/04/2017 13:52	Panorama	male	20	neutral
84	07/04/2017 13:52	Panorama	male	24	neutral
85	07/04/2017 13:53	Panorama	male	35	neutral
86	07/04/2017 13:53	Panorama	male	40	neutral
87	07/04/2017 13:54	Panorama	male	31	neutral
88	07/04/2017 13:54	Panorama	male	37	neutral
89	07/04/2017 13:56	Panorama	male	35	neutral
90	07/04/2017 13:56	Panorama	male	41	neutral
91	07/04/2017 13:57	Panorama	male	36	neutral
92	07/04/2017 13:58	Panorama	male	46	neutral
93	07/04/2017 14:01	Panorama	male	41	neutral
94	07/04/2017 14:02	Panorama	male	38	happiness
95	07/04/2017 14:02	Panorama	female	8	happiness
96	07/04/2017 14:03	Panorama	female	9	happiness

This basic data was used to create summary statistics describing the age distribution of visitors to the painting.

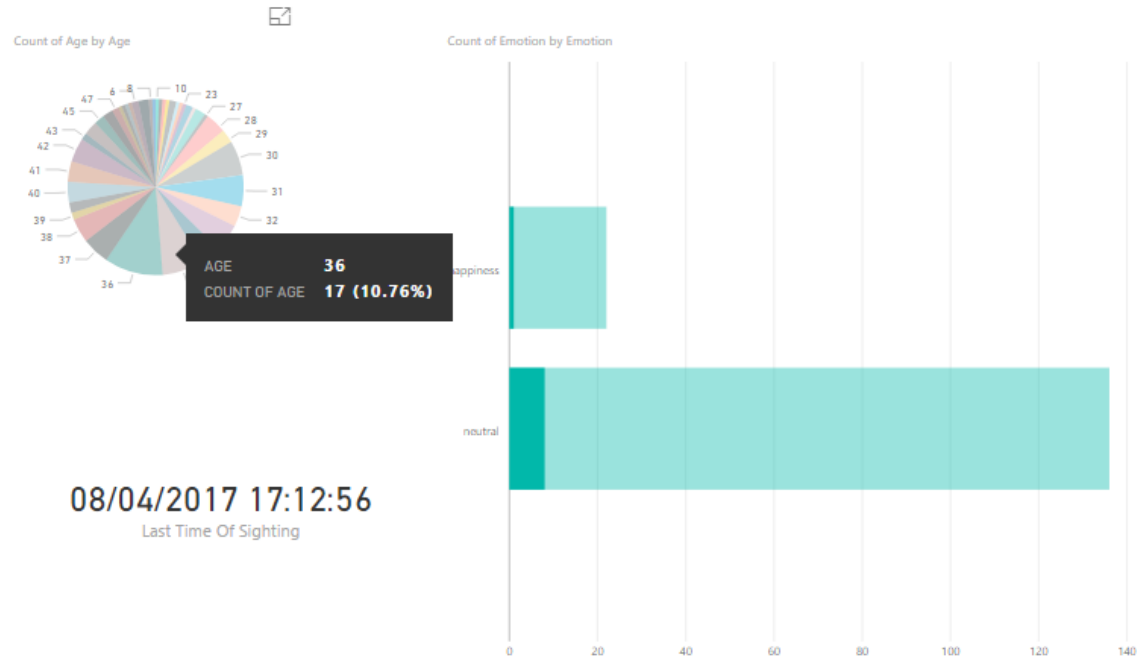

2.3 Privacy

There are naturally some concerns over privacy, however the device makes no record of the images of visitors. No images are stored so it is not possible to review the records of who visited the paint. It is only possible to determine the estimated age and gender of visitors.

3 Further Possibilities

This "proof of concept" was only intended to demonstrate that recognising and responding to visitors was possible. This does raise the potential for further possibilities.

3.1 Recognising Repeat Visitors

No advantage was taken of the devices ability to recognise returning visitors. It would be possible to recognise a visitor returning to the painting and begin the narration where they had left off previously. If the visitor has already heard the complete narration it would be possible to remain silent or deliver a different narration.

The device can be trained to identify particular individuals using a set of about 10 training photographs. So the device could specifically ignore particular individuals, like staff or volunteers. Thus if a member of staff joined a visitor the device would respond to the visitor and respond to the presence of the member of staff. Thus any information about the exhibit is not biased by also including staff and volunteers.

The device can also be trained to recognise particular individual visitors. For example visually impaired visitors (who elected to) could be recognised by the device and be provided with an audio description in addition to the conventional narration.

3.2 Exhibits 2.0

Typically museum and gallery audio visual presentations are relatively fixed and emphasis high production values. A device such as this brings audio visual presentations into the Web 2.0 era, allowing museum and gallery curators, staff and volunteers to create, change and adapt content in response to visitor behaviour.

3.3 Visitor Information

During the project there was not sufficient time to collect any information about how long visitors spent looking at the painting or how the narration affected behaviour. Gathering information about visitor types and dwell time would allow for refinement

of exhibits. Even without the narration, it would be possible to determine how changes in the presentation and labelling of the painting affect behaviour.

With a number of devices installed across a museum it would be possible to determine which exhibits are hotspots and are working best and also which exhibits might benefit from some attention.

4 Next Steps

This "proof of concept" prototype was only ever intended as a "one off" to demonstrate the possibilities provided by cheap and readily available image recognition. Before it could see wider adoption there are a number of immediate next steps.

4.1 Comparison Information

As was mentioned there was not enough time during the project to gathering any information about visitor's behaviour towards the painting. Before audio narration is added to exhibits it would be wise to determine how visitors interact with exhibits without narration. To this end the prototype will be deployed at Shrewsbury Museum (in a better looking box) for 14 days to gather information without the narration running. The information gathered can then be used to provide a summary describing the types of visitors who viewed the painting and their dwell times. This information can then be used to compare the effect of having the narration running.

4.2 Curator's or Non-technical Interface

The device used off the shelf components and Open Source software. However, it does require some technical expertise to configure the device and to update the narration. If the devices it to see wider use it will need an interface that allows curators and other subject matter experts to update the sound files of the narrative. This could be a web interface allowing curators to simply upload new sound files much as they would for social media and sharing websites.

4.3 Multiple Devices

For speed and convenience the prototype was constructed as a single stand-alone device. However the low cost of the device would make it possible to have multiple devices, which either interact or track visitors through a museum or gallery. In order for this to happen the software will need to be adjusted to use different narration for different devices. Currently the design assumes that there is only one narration.

4.4 Information Access and Sharing

Some attempt was made during the project to gather information about how visitors were responding to the painting. To be useful to museum and gallery curators, this information should be presented in a form that suits the museum or gallery professionals needs. The museum and galleries community has a long history of sharing benchmarking and performance information so such an interface could reasonably include comparison information so that the performance of exhibits can be compared across different museums and galleries.

5 Thanks

The team consisted of:

- Joe Collins (Black Radley)
- Ryan O'Neill (Black Radley)
- Richard Wilde (Black Radley)
- Martin Kearn (Microsoft)
- Martin Beeby (Microsoft)
- Paul Foster (Microsoft)

We would like to thank Phil Scoggins and Adrian Perks, for accommodating and supporting the team whilst at Shrewsbury Museum and Art Gallery.