Institution: University of Southampton		
Unit of Assessment: 11 Computer Science and Informatics		
Title of case study: 11-06 Novel image capture, processing and analysis methods for widespread web publishing, new commercial products and cultural heritage preservation		
Period when the underpinning research was undertaken: 2002 – 2020		
Details of staff conducting the underpinning research from the submitting unit:		
Name(s):	**Role(s) (e.g. job title):**	**Period(s) employed by submitting HEI:**
Professor Kirk Martinez	Professor	January 1997 – present
Dr Jonathon Hare	Associate Professor	October 2005 – present
Professor Paul Lewis	Professor	April 1985 – December 2013
Professor Dame Wendy Hall	Regius Professor	January 1984 – present
Dr Philip Basford	Senior Research Fellow	July 2013 – present
Dr David Dupplaw	Research Fellow	June 2003 – December 2013
Dr Sina Samangooei	Research Fellow	July 2009 – June 2014
Dr Simon Goodall	Research Fellow	June 2006 – September 2007
Dr Graeme Bragg	Research Fellow	March 2017 – present
Dr David Young	Research Fellow	April 2015 – April 2016
Period when the claimed impact occurred: August 2013 – December 2020		
Is this case study continued from a case study submitted in 2014? N		

1. Summary of the impact

Software and hardware developed through University of Southampton (UoS) research into the capture, processing and analysis of visual and multimedia data have had far-reaching impacts on web publishing, new machine learning data products and cultural heritage.

- Highly efficient, open source image processing and analysis software libraries have transformed image handling for millions of web users and thousands of companies who benefit from much lower cloud resource requirements, including high-profile users like Wikipedia, Booking.com, the New York Times and Airbnb.
- New commercial products have enabled the Ordnance Survey to deliver on a multi-million-pound contract with the UK Government and explore new markets, and, via a UK start-up, saved analytics firms and banks hundreds of thousands of costly work hours each year.
- Twelve cultural heritage institutions are using UoS custom-built imaging hardware systems, highlighted as best practice by Historic England, to capture images of their collections, including tablets of cuneiform script dating back to 3400 BC, in unprecedented detail.

2. Underpinning research

Interlinked research at the School for Electronics and Computer Science (ECS) has sought to optimise image processing and analysis algorithms and the design of imaging hardware. The group has developed novel, open source software and hardware applications that facilitate faster and more efficient capture, processing and analysis of visual and multimedia data.

Open source image processing and analysis software

Beginning in 2000, ECS had a central role in a series of EU-funded projects that developed new methods for the digitalisation of cultural heritage [**G1**, **G2**]. Collaborators included the National Gallery and the V&A Museum, who were interested in new image capture, content analysis, categorisation and retrieval techniques to better understand and archive their collections, and to widen public access to them. ECS released the software and algorithms developed through these projects into two open-source image processing tools: **libVIPS** and **OpenIMAJ**.

Martinez co-invented libVIPS and optimised it through **G1**/**G2** to process gigabyte-sized multispectral images of artworks. libVIPS is a fast image processing library with low memory needs and is now commonly used as an image pre-processor/resizer to minimise cloud resource usage. It is fully demand-driven; it does not process entire images in memory but instead

streams images through a computer as a series of small regions, which reduces memory use. It was published for the web community on github in 2007 [**3.1**] and ongoing research at ECS [e.g. **G3**] has further optimised its functionality to speed up capture workflow. As a demonstration of its value to the web and research community, the libVIPS library is regularly 'fuzz tested' by Google, a technique for uncovering programming errors in software that could compromise vital web infrastructure, and is part of the PARSEC benchmark suite, which helps to characterise CPU architectures and is widely used by Princeton, MIT and University of Cambridge.

OpenIMAJ [**3.2**] is a broad set of data mining software tools for multimedia content analysis and content generation. Released in 2011 (when it won the ACM Multimedia Open Source Software Competition), it contains implementations analytics techniques that ECS researched and developed via **G1**, **G2** and **G3**. This includes advanced methods for specific tasks like feature extraction through to novel methods for clustering and indexing these features. For example, Hare led the development of a new methodology to combine items from social media streams, such as Flickr photos and Twitter tweets, into meaningful groups that can help users better contextualise and digest the huge volume of information being published [**3.3**]. He developed a method of scaling up a standard image analysis and indexing tool using a distributed cluster of machines [**3.4**]. Many of the algorithms in OpenIMAJ have been successful in international benchmarking contests where algorithms from different research groups and industry are tested.

Novel image analysis techniques

In 2015 the Ordnance Survey (OS) funded Hare and colleagues to develop deep-learning algorithms that mine high-resolution aerial imagery to reveal new information previously locked away in the OS' vast archives [**3.5**]. The research produced image analysis models that, for example, can predict the materials used for roof construction across the country, or facilitate the discovery of new archaeological sites. In 2018, ECS partnered with machine learning firm Evolution AI on Innovate UK project TranscribeAI [**G4**]. The aim was to develop state-of-the-art tools for document analysis that can extract information from semi-structured business documents (e.g. contracts, invoices, accounts). The system developed by Hare and colleagues is able to accurately locate words and numbers in scanned document images to a higher degree of accuracy than any known system, allowing Evolution AI to automate processes often performed manually.

Hardware imaging systems

ECS research into Reflectance Transformation Imaging (RTI), via **G2** and **G3**, led to the development of optimised image capture solutions for cultural heritage [**3.6**]. RTI is a computational photographic method that captures a subject's shape and colour from multiple light directions to reveal surface information that would go unseen under standard examination. Led by Martinez, ECS developed a dome-shaped design – RTI Dome – for higher resolution and faster RTI imaging of ancient artefacts; it is now the world's leading producer of dome-based systems. The group also developed a new structured-light imaging system for capturing images of ancient cylinder seals [**3.7**] – engraved objects that were designed to be rolled into clay to leave impressions.

3. References to the research

3.1 The libVIPS image processing library (published on github in 2007). https://libvips.github.io/libvips/; https://github.com/libvips/libvips

3.2 The OpenIMAJ multimedia content analysis and content generation libraries (published 2011). http://openimaj.org

3.3 J. Hare, S. Samangooei, M. Niranjan, and N. Gibbins. Detection of social events in streams of social multimedia. International Journal of Multimedia Information Retrieval, 4(4):289–302, August 2015. https://doi.org/10.1007/s13735-015-0085-0

3.4 J. S. Hare, S. Samangooei, and P. H. Lewis. Practical scalable image analysis and indexing using Hadoop. Multimedia Tools and Applications, 71(3):1215–1248, Nov. 2012. https://doi.org/10.1007/s11042-012-1256-0

3.5 Zhang, Ce, Sargent, Isabel, Pan, Xin, Li, Huapeng, Gardiner, Andy, Hare, Jonathon and Atkinson, Peter M. (2019) Joint deep learning for land cover and land use classification. Remote Sensing of Environment, 221, 173-187. https://doi.org/10.1016/j.rse.2018.11.014

3.6 Earl, G., Martinez, K. and Malzbender, T. (2010) Archaeological applications of polynomial texture mapping: analysis, conservation and representation. Journal of Archaeological Science, 37. https://doi.org/10.1016/j.jas.2010.03.009

3.7 Martinez, K., Hare, J., Dahl, J., Kelley, K., Young, D. (2018) A structured light approach to imaging ancient Near Eastern cylinder seals: how efficient 3D imaging may facilitate corpuswide research. In, Kelley, Kaye and Wood, Rachel (eds.) *Digital Imaging of Artefacts: Developments in Methods and Aims.* (Access Archaeology) Archaeopress / British Archaeological Reports, pp. 49-74. https://eprints.soton.ac.uk/426431

Key underpinning grants

G1 Three EU/IST projects. 1. An Integrated Art Analysis and Navigation Environment (ARTISTE), EUR2.8m, 2000-2002; 2. Semantic and Content-based multimedia exploitation for European benefit (SCULPTEUR), EUR3.0m, 2002-2005; 3. eCHASE: sustainable exploitation of electronic cultural heritage, EUR3.5m, 2005-2007.

G2 Two EU-FP7 projects. 1. LivingKnowledge: Facts, opinions and bias in time, EUR4.8m, 2009-2012; 2. ARCOMEM: Archive Communities Memories, EUR6.0m, 2011-2013.

G3 AHRC, Reflection Transformation Imaging Systems for Ancient Documents (RTISAD), GBP70,000, 2013.

G4 Innovate UK, Transcribe AI: Artificial Intelligence-Driven Information Extraction from Documents, GBP664,268, 2018-2020.

4. Details of the impact

Novel image capture, processing and analysis tools and methodologies developed and optimised through ECS research have had a far-reaching impact on the way multimedia content is published and developed on the Web, the development of innovative 'big data' products and services for commercial exploitation, and the management of cultural heritage.

Impact of libVIPS/OpenIMAJ libraries on web publishing, software development and teaching

The libVIPS library is now a standard open source Linux package, which has been refined and optimised by an international web developer community and the ECS group itself over the impact period. Recognised as the fastest available image processing library, it is widely used to develop websites and systems that can handle very large numbers of images [**5.1**]. Use of libVIPS minimises cloud resource usage; a saving of >50% of CPU/RAM for each burst of image activity reduces server loads, delivers a more efficient service and negates the need for costly upgrades. The libVIPS source code repository on github has an average of 800 regular users and 3,000 unique visitors each month. But the full scale of its impact is realised through the use of libVIPS as a core driver of user-facing web applications and software tools. In 2014, libVIPS was used to create Sharp, a high-performance module for resizing and formatting large images to smaller, web-friendly JPEG/PNG images of varying dimensions [**5.1**]. Sharp acts as a 'web wrapper' in that it provides access to libVIPS for web developers by making it accessible to javascript. By the end of 2020, Sharp had reached 1,800,000 downloads per week [**5.1**], while 640 websites are known to use libVIPS via Gatsby, a content management system that helps developers build ultra-fast websites [**5.1**]. While it is not possible to quantify the full reach of libVIPS, the founder of Sharp wrote [**5.1**]: '*Over 23,000 open source projects rely on sharp, including most JavaScript-based content management systems. This means you've almost certainly visited a website in the last week that used libvips*.'

The widespread use of libVIPS is demonstrated through end user case studies. In September 2013, *VipsScaler*, a wrapper to libVIPS, was rolled out across all Wikimedia and Wikipedia sites [**5.2**]. This allowed, for the first time, the creation of thumbnails of very large PNG files due to significant reductions in required memory [**5.2**]. Amazon Web Services (AWS) uses libVIPS for efficient image resizing within one its main applications, Lambda@Edge – a tool that allows developers to run code in the AWS code without provisioning or managing servers [**5.3**]. This has a significant knock-on impact. For example, libVIPS was used, via AWS services, by the New York Times to overhaul its digital archive and create, store and serve billions of small images for subscribers interested in back issues of the newspaper dating back to 1851 [**5.4**]. Other high-profile users of the faster, more comprehensive image handling offered by libVIPS include Booking.com, AirBnB Engineering & Data Science, the US Food and Drug

Administration, the bank Capital One, Braun, Pantene and US and Spanish property listing sites HomeFinder and Idealista [**5.1**, **5.5**].

HomeFinder reports that libvips '*is fast enough that we can handle all the photo (real estate listings) manipulation on-demand, giving us the flexibility to use any photo size we want on the site, while only processing photos that are actually seen by our visitors*' [**5.1**]. Swiss digital agency Liip uses libVIPS to generate 100,000 images per day for clients that include Migros, Switzerland's largest retailer and the country's largest private employer; libVIPS has cut Liip's response time for rendering images by three to four times and has saved them money by removing the need to scale up their cloud servers [**5.1**]. Slovenian company Hooray Studios uses libVIPS for their personalised children's books service to process orders faster; the company doubled its revenues in 2019 to EUR27,000,000 and now has subsidiaries in the US, UK, Italy and Germany [**5.1**]. As project partners on the European grants [**G1**, **G2**], the The National Gallery London and The Louvre uses libVIPS. Others including National Gallery of Washington use it to convert their collection to multi-resolution files [**5.1**].

The OpenIMAJ libraries are used by lecturers and web developers across the world, including large companies like Yahoo and Comcast (the former CTO of Comcast contributed to the development of OpenIMAJ through **G2**) [**5.6**]. They allow developers to quickly create applications that extract information from multimedia data. Around 31,000 code files mention OpenIMAJ on GitHub and it features in 1,461 'Maven pom files', which represent individual software projects by users of the library [**5.6**]. It is a key component of very widely used software, for example the Webcam Capture library that allows webcams to be accessed directly from Java [**5.6**]. OpenIMAJ is used in university teaching internationally; for example, the University of Missouri uses OpenIMAJ in a module on big data analytics [**5.6**]. It is also used as the basis for the code in a number of textbooks on data visualisation techniques, Raspberry Pi and feature extraction [**5.6**].

Impact of new commercial data products on businesses

The collaboration with Ordnance Survey [**3.5**] on new ways to exploit its aerial imagery archives using machine learning has facilitated business innovation in the government-owned company. Before the partnership, methods for creating OS's data products were largely manual and therefore limited in scope due to the high costs involved. Noting that '*commercial success*' would not have occurred without the ECS research contribution, OS cites the key impacts as [**5.7**]:

- Enabling the OS to '*efficiently meet some of the challenging new data needs*' of a multi-million-pound contract with the UK Government to deliver the Public Sector Geospatial Agreement – a commitment to opening up access to geospatial data to support the delivery of critical infrastructure and services.
- Opening up new multi-million-pound markets.
- Identifying opportunities for making up to GBP5,000,000 in annual savings.
- Granting of several patents in the UK, Europe and US.

The document analysis system created for Evolution AI [**G4**], which intelligently extracts data from corporate and financial documents, is the foundation of all of the company's commercial products [**5.8**]. It allowed Evolution AI, for the first time, to automate processes that were previously undertaken manually, bettering human accuracy and reducing the time taken (and therefore cost) '*from days to seconds*' [**5.8**]. The impact is demonstrated through client case studies [**5.8**]:

- Saved US analytics firm Dun & Bradstreet 50,000 hours of manual work a year (28 FTEs).
- Saved RBS (NatWest Group) between 100,000 and 200,000 hours of manual work a year and reduced their invoice processing times from 24 hours to three seconds.
- UK asset management firm Unigestion achieved 85% cost savings, increased data accuracy from 30% to 99.5% and halved their data extraction time.
- Made data extraction processes for US analytics firm Fitch Solutions 15 times faster.

Impact of imaging hardware systems on the cultural heritage sector

Fifteen custom-built RTI systems, developed through ECS research [**3.6**, **3.7**, **G2**, **G3**], were supplied to 12 leading museums and universities over the impact period [**5.9**]. Institutions use these systems on a regular basis to capture images of their artefacts and better understand their collections, which in turn improves the quality of their public engagement. The systems have been used extensively to image cuneiform script; the oldest form of writing in the world (first used in around 3400 BC) distinguished by wedge-shaped marks on clay tablets [**5.9**]. Thousands of tablets have been captured with the RTI domes and logged with the global repository at UCLA: the Cuneiform Digital Library Initiative [**5.9**]. Historic England featured Southampton's RTI dome imaging techniques as best practice in its guidelines for Multi-light Imaging (2018) [**5.10**].

Figure 1: A Cuneiform tablet (left) and an RTI Dome system (right)

Examples of impact on cultural heritage institutions during the impact period include [**5.9**]:

- Since 2014 the British Museum has used their Dome to image objects of stone, paper, jade and metal, including 15th-17th Century metalpoint drawings for a 2015 exhibition.
- The Louvre (from 2015) is using their Dome to image its 20,000 cuneiform tablets.
- The Ashmolean Museum (from 2014) in Oxford is using the Dome to image a wide range of objects, including its collection of 5,000-6,000 cuneiform tablets.
- The Direction de l'Archéologie du Pas-de-Calais (from 2017) uses it to photograph coins, lithics and ceramics; it enables improved technical drawings of the objects.
- The National Museum of Antiquities, Netherlands (2020) uses theirs to image objects.

5. Sources to corroborate the impact

5.1 Evaluation report of the impact of libVIPS: statistics and case studies from end users. Corroborating contact: Author of Sharp
5.2 Wikimedia public release on the implementation of VipsScaler (built around libVIPS) to improve image handling on Wikimedia and Wikipedia sites, September 12, 2013. https://diff.wikimedia.org/2013/09/12/vipsscaler-implementation-wikimedia-sites/
5.3 Blog post by Amazon Web Services on resizing web images with its Lambda@Edge application (citing the use of libVIPS via Sharp to do so), February 20, 2018. https://aws.amazon.com/blogs/networking-and-content-delivery/resizing-images-with-amazon-cloudfront-lambdaedge-aws-cdn-blog/
5.4 New York Times developer on their use of libVIPS via Sharp in AWS, May 27, 2014: https://github.com/lovell/sharp/issues/35#issuecomment-44336416
5.5 libVIPS Wikipedia page specifying key users:https://en.wikipedia.org/wiki/VIPS_(software)
5.6 Evaluation report of the impact of OpenIMAJ: statistics and examples of end users.
5.7 Corroborating statement from Director of Propositions and Innovation, Ordnance Survey.
5.8 Corroborating statement from the Chief Technical Officer, Evolution AI; Client case studies via the Evolution AI website: https://www.evolution.ai/case-studies
5.9 Evaluation report of the impact of ECS RTI systems: corroborating emails and news releases. Corroborating contact: Professor of Assyriology, University of Oxford