Institution: Imperial College London		
Unit of Assessment: 11 Computer Science and Informatics		
Title of Case Study: SLAM: Enabling Algorithms for Robotics and Augmented Reality		
Period when the underpinning research was undertaken: 2005 – 2019		
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:**
Andrew Davison	Professor	2005 – present
Stefan Leutenegger	Senior Lecturer	2014 – present
Paul Kelly	Professor	1989 – present
Richard Newcombe	PhD Student	2008 – 2012
Renato Salas-Moreno	PhD Student	2012 – 2016
John McCormac	PhD Student	2014 – 2018
Period when the claimed impact occurred: 1 Aug 2013 – 31 Dec 2020		
Is this case study continued from a case study submitted in 2014? No		

1. Summary of the impact

Vision-based Simultaneous Localisation and Mapping (SLAM) is a fundamental and broadly applicable technology in AI that enables devices to operate intelligently by mapping the general unstructured world using low-cost cameras and efficient on-board processing. The Imperial team have made a sequence of highly influential advances in SLAM algorithms: (i) sparse visual feature-based mapping enables drift-free long-term 3D localisation; (ii) dense scene mapping provides detailed scene modelling; and (iii) semantic mapping brings object understanding into maps. These new algorithms have been harnessed in worldwide consumer products as key enabling features of Dyson's first ever robotic products, the Dyson 360 Eye and Heurist robot vacuum cleaners. SLAM is also used in augmented reality headsets to track and map surroundings, and Imperial's algorithms have contributed to both Microsoft's ongoing Kinect and Hololens products as well as systems at Facebook/Oculus via the acquisition of Imperial's start-up Surreal Vision for USD40,000,000. A second Imperial spinout, SLAMcore, has received over GBP7,000,000 of investment to date, employs >20 people, and commercialises Imperial's SLAM software for broader robotics applications.

2. Underpinning research

Research progress in vision-based SLAM can be understood in terms of interacting layers of capability in a software stack for vision applications, and Davison and Leutenegger have worked for over 15 years at the forefront of research across this stack. Their research work has enabled low-cost real-time spatial intelligence for a new generation of consumer devices.

Davison's seminal work on the MonoSLAM algorithm [R1] was the first demonstration of real-time SLAM with a single camera. Incremental probabilistic filtering for scene mapping had previously been shown in robotics only using specialised and expensive sensors, such as laser range-finders. MonoSLAM's innovation in sparse active feature search, 3D dynamics modelling and sampling-based landmark initialisation proved that probabilistic 3D SLAM was possible with a low-cost standard camera, opening up new commercial applications in mass-market products. MonoSLAM builds a persistent 3D sparse landmark map of an unknown environment in real-time and uses it to enable drift-free long-term localisation. Davison's research results in Inverse Depth Feature Representation [R2] improved feature initialisation, especially in large or outdoor scenes. These methods are now at the heart of the SLAM system in Dyson's 360 Eye robotic product and inspired other companies to develop visual SLAM systems for robotics. Sparse visual SLAM for device tracking is also critical in the exploding field of augmented reality (AR) on smartphones or mobile headsets. MonoSLAM initiated a huge research interest and commercial focus in this area, as indicated by the high number of citations to papers from the Imperial team and the investment of companies such as Dyson and Facebook.

The next level in SLAM is dense scene mapping, which uses extra processing and diverse camera hardware to capture detailed 3D scene information for planning, obstacle avoidance and scene augmentation. The Imperial team's collaboration with KinectFusion [R3] and their work on DTAM [R4] resulted in pioneering algorithms using parallelisable optimisation methods that run on GPU hardware to enable real-time dense scene mapping for the first time. KinectFusion was the first SLAM algorithm that was integrated with Microsoft's new Kinect consumer-level depth camera for robotics and AR. DTAM showed that detailed real-time scene reconstruction can be achieved even with a regular single camera and that camera tracking can be robust without features via whole image alignment. Dense SLAM for AR has been commercialised by the Imperial team through the co-founding of Surreal Vision, which was subsequently acquired by Oculus/Facebook.

Finally, the Imperial team's work on SLAM++ [R5] and SemanticFusion [R6] led to key advances with respect to the challenge of adding a semantic layer to SLAM so that robots and other devices can understand the identities of scene entities as well as their geometry. The research on SLAM++ showed that real-time 3D object detection allows an "instance-level" graph of key objects in a scene to be built and optimised – the objects themselves becoming high-quality sparse landmarks. For the first time, SemanticFusion enabled the output from a convolutional neural network (CNN) trained for image segmentation to be fused incrementally into a dense SLAM map that supports full pixel-wise semantic labelling of a scene. Semantic scene understanding has proven crucial for applications that require intelligent interaction between a robot and its environment (e.g., for manipulation, or task and motion planning) and between a user and a robot (e.g., via natural language queries).

3. References to the research

[R1] MonoSLAM: Real-Time Single Camera SLAM. A. J. Davison, I. Reid, N. Molton and O. Stasse, IEEE Trans. Patten Analysis and Machine Intelligence (PAMI) June 2007, vol. 29, pp. 1952 – 1067. 3996 citations, Google Scholar. (Collaboration with University of Oxford.) https://www.doc.ic.ac.uk/~ajd/Publications/davison_etal_pami2007.pdf

[R2] Inverse Depth Parametrization for Monocular SLAM. J. M. M. Montiel, J. Civera and A. J. Davison, IEEE Trans. on Robotics Oct. 2008, 24(5), pp. 932 – 945. 844 citations, Google Scholar. (Collaboration with University of Zaragoza, Spain.) https://www.doc.ic.ac.uk/~ajd/Publications/civera_etal_tro2008.pdf

[R3] KinectFusion: Real-Time Dense Surface Mapping and Tracking. R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohli, J. Shotton, S. Hodges, and A. Fitzgibbon, Int. Symp. on Mixed and Augmented Reality (ISMAR) 2011, pp. 127 – 136. 3755 citations, Google Scholar (Collaboration with Microsoft Research, Cambridge.) https://www.doc.ic.ac.uk/~ajd/Publications/newcombe_etal_ismar2011.pdf

[R4] DTAM: Dense Tracking and Mapping in Real-Time. R. A. Newcombe, S. Lovegrove and A. J. Davison. Proc. IEEE ICCV, Nov. 2011. 1808 citations, Google Scholar. https://www.doc.ic.ac.uk/~ajd/Publications/newcombe_etal_iccv2011.pdf

[R5] SLAM++: Simultaneous Localisation and Mapping at the Level of Objects. R. F. Salas-Moreno, R. A. Newcombe, H. Strasdat, P. H. J. Kelly and A. J. Davison, IEEE Proc. CVPR, June 2013. 683 citations, Google Scholar. https://bit.ly/3rEQXFN

[R6] SemanticFusion: Dense 3D Semantic Mapping with Convolutional Neural Networks. J. McCormac, A. Handa, A. J. Davison, S. Leutenegger. Proc. IEEE ICRA, May 2017. 325 citations, Google Scholar. https://bit.ly/36V5t4x

4. Details of the impact

Impact on robotics products and R&D investment

Dyson product development: Davison worked as an academic consultant to Dyson until 2014 on the long-term project of developing their **first robot product (Dyson 360 Eye)**, which makes direct use of the published algorithms in [R1, R2], and incorporates them into Dyson's robotic vacuum cleaner development. He was responsible for guiding Dyson engineers to successfully design an omnidirectional visual SLAM system, which was the unique key feature (and thus chosen as the name) of the Dyson 360 Eye released worldwide in 2015. Dyson is one of few large UK companies to have a leading robotic product already on the market.

In 2018, Dyson's second robotic product, the **Dyson 360 Heurist** [I2], which also incorporates Davison's core SLAM system, went on sale worldwide. The Principal Robotics Architect at Dyson states "*Andrew's pioneering research on vision-based SLAM in published algorithms such as MonoSLAM and Inverse Depth Features directly inspired the visual mapping system Dyson developed for the robot, based on a 360-degree panoramic camera. This enabled the robot to navigate precisely and repeatedly, and therefore clean rooms in a much more direct and efficient manner than other products on the market […] The core visual SLAM system developed for the 360 Eye with Andrew's input lives on in Dyson's latest and current robot vacuum cleaner, the Dyson 360 Heurist, which is now on sale worldwide"* [I1].

In addition, Imperial's SLAM research led to Dyson making large ongoing investments in robotics as a core focus of their business. At the start of Davison's consultancy, the internal Dyson robotics team had fewer than 5 people, and now **Dyson employs hundreds of scientists and engineers in robotics R&D** in both the UK and Singapore [I1]. In 2020, Dyson announced a new investment of GBP3 billion specifically into AI, robotics and batteries [I9]. Since 2014, the Imperial team continues to collaborate closely with Dyson, who funded the Dyson Robotics Laboratory at Imperial directed by Davison and Leutenegger [I3]. The lab has produced over 10 patents for technologies selected by Dyson as relevant to their future robotics product lines.

Impact on augmented reality products and R&D investment

The Imperial research on sparse, dense and semantic SLAM has also directly impacted the augmented reality industry, specifically VR and AR products in headsets and smartphones, which need accurate but low-cost tracking and scene reconstruction from their on-board cameras. Major multi-billion dollar industry projects such as **Facebook/Oculus' VR/AR platforms, Microsoft Hololens, and Google's ARCore** have acknowledged that their algorithms used for sparse and dense mapping are derived from Imperial research work.

Google initiated Project Tango in 2013 to put SLAM into smartphones. In 2018, this technology was transitioned into ARCore, a major platform which enables developers to build AR applications for Android, now used in over 250 million devices worldwide, for functionality such as the AR walking directions feature of Google Mobile Maps. The Engineering Director at Google states [I11]: "*Imperial College's visual SLAM research led by Andrew Davison and Stefan Leutenegger was one of the initial inspirations behind Project Tango, going back to the first demonstration of real-time single camera MonoSLAM which showed the feasibility of vision-driven localisation and mapping using low-cost sensors. […] As the market for spatial aware computing evolves, research from Imperial College on SLAM and semantic understanding of the environment will continue to grow in demand. In my opinion, the impact of the technology will be similar to that of the Global Positioning System (GPS) which has transformed many existing* *industries, and made new ones possible. Wearables, smart home devices, vehicles, and robotics will all depend heavily on a shared understanding of physical space.*"

Surreal Vision acquisition: In 2013, Imperial PhD students (Newcombe, Lovegrove and Salas-Moreno) co-founded the AR company Surreal Vision, advised by Davison and Kelly, to commercialise the dense and object-based SLAM research in DTAM [R4] and SLAM++ [R5]. Based on this work, Surreal Vision developed a new level of accurate but efficient 3D reconstruction algorithms for VR and AR systems, able to provide a rich model of an environment including people and their interactions, but still running on modest consumer computing hardware. Surreal Vision was acquired for USD40,000,000 by Oculus (part of Facebook) in 2015 [I4]. The Surreal Vision team is now part of Facebook Reality Labs, a world leading industrial research team in AR for telepresence. Newcombe is Director of Research Science and leads a team of over 100 researchers which created the fundamental LiveMaps SLAM technology that originates from Imperial's research. This is a key part of Facebook's future technology plans for AR as a cloud platform for building, sharing and annotating always up-to-date scene maps across a multitude of user devices [I8].

Microsoft product development: Imperial's SLAM work was incorporated in Microsoft's Hololens AR headset product line through multiple stages. MonoSLAM [R1] directly impacted the development of Hololens' core sparse visual tracking system [I5]. After that, in 2010, KinectFusion [R3], a collaboration between the Imperial team and Microsoft Research Cambridge, became a key Microsoft technology and is now part of Microsoft's official Kinect for Windows SDK, which is used in millions of Kinect Sensors [I6]. KinectFusion also enables the dense mapping capabilities of Microsoft's Hololens 1 (on sale in 2016) and Hololens 2 (released in 2019 and now on sale) products, which feature Kinect-like depth sensors [I5].

The Partner Director of Research, Microsoft Research Cambridge states "*[t]he SLAM research at Imperial College led by Andrew Davison has impacted the development of HoloLens significantly. Andrew's research on real-time MonoSLAM was an early inspiration in the development of the core tracking system in HoloLens. HoloLens used a sparse feature mapping visual SLAM approach to achieve the robust and accurate inside-out tracking of the wearer's head which is crucial for a convincing user experience of mixed reality. […] KinectFusion was very important in impacting the approach we currently use for environment mapping in HoloLens 1 and 2. KinectFusion pioneered real-time volumetric fusion of depth data, and the way that this can be implemented efficiently using parallel computation. […] the KinectFusion algorithm for dense environment mapping based on measurements from a moving depth camera […], originally targeted Microsoft's newly release Kinect device, and was hugely influential guiding research in scene mapping for multiple applications including robotics, with the unprecedented detail and speed of reconstruction it enabled. […] On the HoloLens, a derivative of the KinectFusion algorithm allows a real-time reconstruction of the environment around the user, allowing virtual objects to be aligned with real world objects and occlusions of the virtual by the real to be handled appropriately*" [I5, I6].

Impact on software development in robotics industry

SLAMcore is an Imperial spin-out company, co-founded in 2016 by Leutenegger, Davison, two PhD students (Zienkiewicz and Kim) and a lab manager (Nicholson). As of 2020, it has over 20 employees and received investment from the UK, US and Japan of over GBP7,000,000 to develop Spatial AI solutions for robots and drones [I7].

SLAMcore has impacted the market of spatial robot and drone software development by offering a licensable software development kit (SDK). The SLAMcore Spatial AI Software Stack SDK enables robots to determine their position and map their surroundings in detail to enable long-term intelligent operation from on-board cameras and other sensors. SLAMcore has established formal hardware partnerships with Intel, NVidia and ARM [10]. The SDK currently is licensed by monthly subscription by 8 customer companies in applications such as warehouse management and evaluated by over 50 more that require SLAM technology for new robot products.

Beneficiaries of the impact

Imperial's vision-based SLAM research has led to key low-cost localisation and mapping algorithms used by major international technology companies, such as Dyson, Facebook, Google and Microsoft, enabling development of new spatially intelligent consumer products and services. Understanding their context in the surrounding environment enables products such as autonomous cleaning robots to relieve users from household cleaning chores and human-centred, mixed-reality headsets to facilitate new forms of entertainment as well as workplace tasks.

Nature of the impact

Vision-based SLAM is an enabling technology for low cost, high performance, mobile robot products emerging in the consumer sector. The research has also directly impacted development and products in the growing area of mixed-reality devices, which global technology companies now regard as a major new wave in future human-computer interaction.

5. Sources to corroborate the impact

[I1] Letter from the Dyson Robotics Systems Architect: Principal Engineer, corroborating the contribution of the Imperial Team in the development of the Dyson 360 Range of Robotic Vacuum Cleaners and in expanding the Dyson Robotics Team.

[I2] Dyson 360 Heurist current product website, 2020. https://bit.ly/2O2DVne Link archived here.

[I3] Dyson Robotics Lab at Imperial College announced, National news, 2014. https://www.bbc.co.uk/news/technology-26084765 (The announcement also appeared on the Sunday Times front page.) Link archived here.

[I4] Imperial SLAM spinout Surreal Vision acquired by Facebook/Oculus announcement, mentioning Imperial College, 2015. https://ocul.us/3kh6qZa Link archived here.

[I5] Letter by the Microsoft Partner Director of Science corroborating the Imperial Team's contribution to Hololens 1 & 2 and Kinect Devices.

[I6] Microsoft Windows SDK for Kinect v2 sensors including KinectFusion, 2014. http://bit.ly/3oIWtWf Link archived here.

[I7] SLAMcore funding round announcement. Sep. 2018. http://bit.ly/2O4KVzM Link archived here.

[I8] Facebook Reality Labs blog where Dr Richard Newcombe discusses the impact of SLAM on AR and VR applications, 2019. https://bit.ly/3kfzy31 Link archived here.

[I9] Dyson to invest GBP3 billion in AI, robotics and batteries, 2020. https://www.ft.com/content/834a568d-5bd1-4fd5-a7d3-c66d2eed9265 Link archived here.

[I10] SLAMcore technology blog, including details about hardware partnerships with Intel, NVidia, ARM. https://blog.slamcore.com Link archived here.

[I11] Letter by the Engineering Director at Google, corroborating the contribution of the Imperial Team to the development of Google Project Tango and ARCore.