25th International Congress on Modelling and Simulation, Darwin, NT, Australia, 9 to 14 July 2023
mssanz.org.au/modsim2023

The importance of self-awareness to digital twins

Lachlan Hetherton, Matt Bolger, Ping Xie, Zikai Zhang and Chris Rucinski

CSIRO Data61, Melbourne, Australia
Email: lachlan.hetherton@csiro.au

Abstract: The growing need for automation, decision support and predictive analytics systems in real world scenarios has given rise to the research field of the "digital twin"; a live, digital replica of a physical object, process, or system. Digital twins provide live "situational awareness" by combining cyber-physical systems, artificial intelligence, computational modelling, computational process orchestration and human-computer interfaces (e.g., mixed-reality visualisation) for the purpose of decentralising, or entirely automating, decision-making processes that relate to a product or service.

A great deal of research into digital twins focusses on either modelling (researching methods or building systems to digitally replicate a physical system), or data acquisition (researching methods or building sensors for more effective data capture). As a research team focussed on translating research into real-world applications, we are interested in the higher-order challenges that arise during the laboratory testing and early operationalisation stages of the development process of digital twins. Solutions to these challenges are critical for any digital twin to be practical in a real-world setting. Common among these challenges is the requirement for a digital twin to observe its own components and capabilities (i.e., the interfaces between itself and the physical system), and to include these observations as part of its situational awareness output for use by decision makers.

In this paper, we explore the challenges that emerge when attempting to implement and maintain a digital twin in a laboratory or commercial / operational setting; particularly where visual sensing is employed as the primary means of sourcing data. We describe our team's approach to incorporating self-awareness capabilities into the Commonwealth Scientific and Industrial Research Organisation (CSIRO)'s Mixed Reality Lab (MRL) technology, and how these capabilities contribute to the ongoing task of applying the MRL technology to a variety of industrial research challenges.

Keywords: *Digital twins, situational awareness, mixed reality, workspace, computational modelling*

1. INTRODUCTION

We define a “digital twin” to be a living, digital replica of a physical object, process, or system, composed of three elements: (1) a physical system, (2) a digital model of that system (the twin), and (3) the information flow between these two systems (Watkins, 2021). Underpinning these three elements is sophisticated information management and computational process orchestration infrastructure that, for the twin to provide value as a truly up to date, “living” replica, ensures timely bidirectional information flow between the physical system and its digital replica. In a practical sense (Figure 1), data streams, sensors or people gather information about the current state of the physical system and regularly inform the digital twin, which assimilates this information and provides a model of “situational awareness”. Situational awareness can be defined as the perception of the physical system within a defined volume of time and/or space (Endsley, M.R., 1995). The utility of that perception depends entirely on the reliability, integrity, and temporal/spatial resolution of the data available to the digital twin.

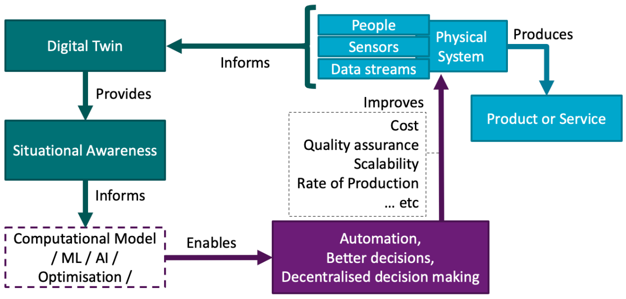

Figure 1. Conceptual model of an ecosystem involving a digital twin

The situational awareness model is relied upon by one or more domain-specific algorithms (e.g., optimisation, machine learning, computational modelling, etc). Finally, outputs of these algorithms are used for task automation, or presented to agents in the physical system to autonomously make system-critical decisions that previously would have required escalation to higher tier decision makers (i.e., decentralised decision making). In the sections below, we will describe our experiences building and maintaining digital twins and explore several higher-order challenges we encountered during laboratory testing.

2. A PLATFORM FOR BUILDING AND MAINTAINING DIGITAL TWINS

There are several ongoing research projects at CSIRO that require the implementation of a situational awareness capability for decision-support, each requiring a continuous, visual assessment of a physical system. To satisfy the common requirements of these projects, the research team developed a software platform that can be configured to meet the individual needs of any given visual assessment-related digital twin. This platform is CSIRO’s Mixed Reality Lab (MRL) software suite (Watkins et al, 2021), built using the Workspace framework (Cleary, 2020). Some examples of projects that employ the facilities of the MRL are as follows:

- Determining the extent to which a large, manufactured object/part deviates from its specification (Hetherton, 2019). The aim is to rapidly identify errors on a large proportion of parts earlier in the production process, avoiding expensive part rectification prior to final part assembly, which, due to global supply chains, can require transport of parts internationally at great cost. It has been shown that accurate measurements can have a 10-15% impact on production costs alone (Czichos 2011).
- Human performance engineering for the prevention of injury in the workplace (Cohen et al, 2017). According to Safe Work Australia, between 2008 and 2018 workplace injuries in Australia reduced productivity by 2.2 million Full-Time Equivalent (FTE)’s (Safe Work Australia, 2022). As automation continues to become a part of the modern workplace, humans and machines are increasingly required to

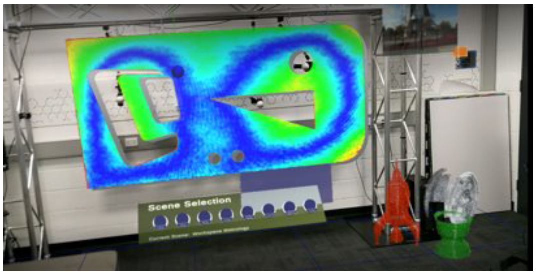

Figure 2. An augmented reality visualisation produced by the Mixed Reality Lab, showing an accuracy overlay on top of a manufactured part

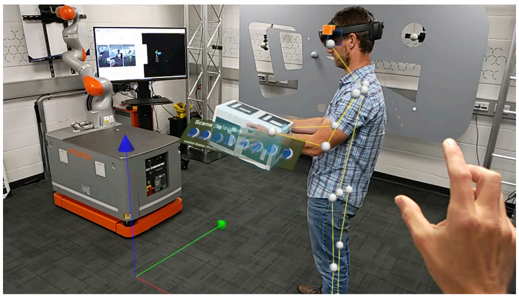

Figure 3. An augmented reality visualisation produced by the Mixed Reality Lab used to demonstrate safe lifting practices on a work site

operate within the same work envelope, increasing the likelihood of injury. We can reduce risk through situational awareness of not only the state of machinery, but of the state of human operators and the risk placed upon them.

The MRL platform permits solutions to be tailored to each specific project but composed from a suite of common capabilities. The capabilities that define the fundamental utility of the MRL and that are shared among the various research projects are as follows:

- **Data capture and storage:** Data must be acquired from sensors, human operators, and external databases / Application Programming Interface (API)'s at regular intervals. In the case of the above listed projects, both commercial and consumer-grade optical and time-of-flight sensors are used.
- **Configurable sensor network:** Sensor network configurations (the types of sensors required, their locations and the underlying topology) vary for each specific application, based on the physical properties of the physical system (scale, material properties, etc); the conditions under which observations are being made (lighting, object is in motion, etc); and the accuracy requirements of the inspection (product specifications, regulatory requirements, etc).
- **Continuous sensor calibration and registration:** The reality of implementing a digital twin is that sensors are frequently disturbed both accidentally (e.g., heating/cooling of fixtures) or deliberately (e.g., sensors adjusted for change in product line). When this occurs, recalibration is required. It is therefore vital that the system provide easy-to-use, efficient calibration tools.
- **Process orchestration and task distribution:** In any given project, process flows and compute networks differ and change over time. The system must facilitate different configurations and process flows. We term an instance of a set of related processes/tasks an "Itinerary".
- **Image processing:** Standard computer vision techniques, such as algorithms in the OpenCV library (Culjak 2012), are required to feed information to downstream feature detection or reconstruction algorithms, or deep learning (DL) algorithms such as OpenPose (Cao et al. 2019) or Detectron2 (Wu et al. 2019) are required to identify body joints and other points of interest in video streams.
- **3D reconstruction**: In many cases that relate to visual assessment of objects/processes, it is necessary to construct a 3D model of the physical system. Frequently, a 3D point-cloud reconstruction is used. At present, Structured Light Systems (Rebo, M., Brandner, M., 2005) and CSIRO's patented Stereo Depth Fusion (SDF) are available (Stainlay, 2005) within the MRL.
- **Analytics and evaluation**: Application-specific workflows are needed to compare the digital reconstruction with reference data, such as Computer Aided Design (CAD) models, of the physical system. These results can then be statistically analysed for automation and decentralised decision making.
- **Visualisation and decision making**: The results of the 3D reconstruction, analytics, and evaluation can be presented to the user and/or participant. This can be as complex as an Augmented Reality (AR) visualisation or as simple as a traffic light indicator showing the operator whether the results are "pass" or "fail".

While in principle these capabilities are all that is required to deliver a digital twin when the system is in complete working order, higher-order problems arise during the laboratory integration test of an application-specific implementation of the MRL technology, which are discussed in detail in section 3.

3. HIGHER-ORDER CHALLENGES

During the laboratory testing phase several projects, the team identified a fundamental data gap present in the initial iterations of the platform – the state of the digital twin system itself, i.e., self-awareness. Without this information, managing the system's deployment and maintaining it over its lifetime would require prohibitive numbers of human resources to manually monitor and debug problems. The specific challenges the team encountered are as follows:

- Repeated, unexplained, intermittent failure of, or loss of connection to, a specific device or machine. Examples encountered during laboratory integration testing included faulty data transfer or power cables, cables being accidentally disconnected, faulty device firmware from vendors, untimely remote device updates pushed by vendors, and lost data due to exceeding available network bandwidth. In each case, senior engineers were required on site for issue diagnosis and resolution, as junior engineers / lab technicians (employed to reduce cost) found diagnosing the cause of system failures frequently outside their limited expertise.
- Difficulty diagnosing the cause of problems in a large, distributed data pipeline. In the MRL, a single data capture can produce in the order of gigabytes of data. If results are not as we expect, we must be able to find and explore the data to diagnose problems – has a sensor moved? has a denoising process

failed? has there been an algorithm change that has negatively affected results? If intermediate results are discarded, this is not possible. Furthermore, manually searching for results on a file system and inspecting them using an external tool is difficult, time consuming, and error prone.

- Frequent unintentional changes to the physical device configuration. This causes the calibration to be in error and the system unable to produce results to the required fidelity. Changes occur for many and varied reasons, ranging from stakeholders touring the site and touching devices, to environmental temperature fluctuations loosening fixings, to birds (yes, birds) landing on top of devices.
- Partial data loss due to file system corruption, hard disk failure or human error. It is not scalable, or reliable, to have engineers manually re-run intermediate processes to produce replacement outputs. Process inputs and intermediate results need to be properly recorded so that outputs can be reproduced without having to manually re-run any intermediate processes. This is particularly important for organisations with business models that rely on demonstrating of high manufacturing quality.

The research team implemented several solutions to these issues, each described in the sections below.

3.1. System status dashboard

To diagnose intermittent failures more quickly and to help identify changes to devices, the team created the web-based *Status Dashboard* application that uses a traffic light system to indicate areas that the operator should pay attention to. The dashboard uses a Representational State Transfer (REST)-ful API for both submitting and querying status reports. This dashboard sits on top of the same database that is used by the core digital twin application suite.

Reports are submitted via HTTP POST requests by daily / hourly Workspace workflows that are scheduled by the central continuous integration server (Jenkins, 2023), and can relate to any given source item or category, allowing the team to expand the types of reports as needed. At present, reports are submitted on behalf of; the database server; each sensor and machine in the data acquisition/processing network; each task itinerary submitted for processing (e.g., any time a 3D reconstruction is requested); and each Jenkins job (e.g., for each integration test run.

In the device network section of the application (Figure 4, left), we can see the sensor network presented on a grid that corresponds to the physical system's coordinate system, with each sensor placed according to its approximate coordinates (obtained in the latest calibration of the selected category), on the selected axes. Any machines or sensors that are missing calibration information are listed at the bottom of the chart.

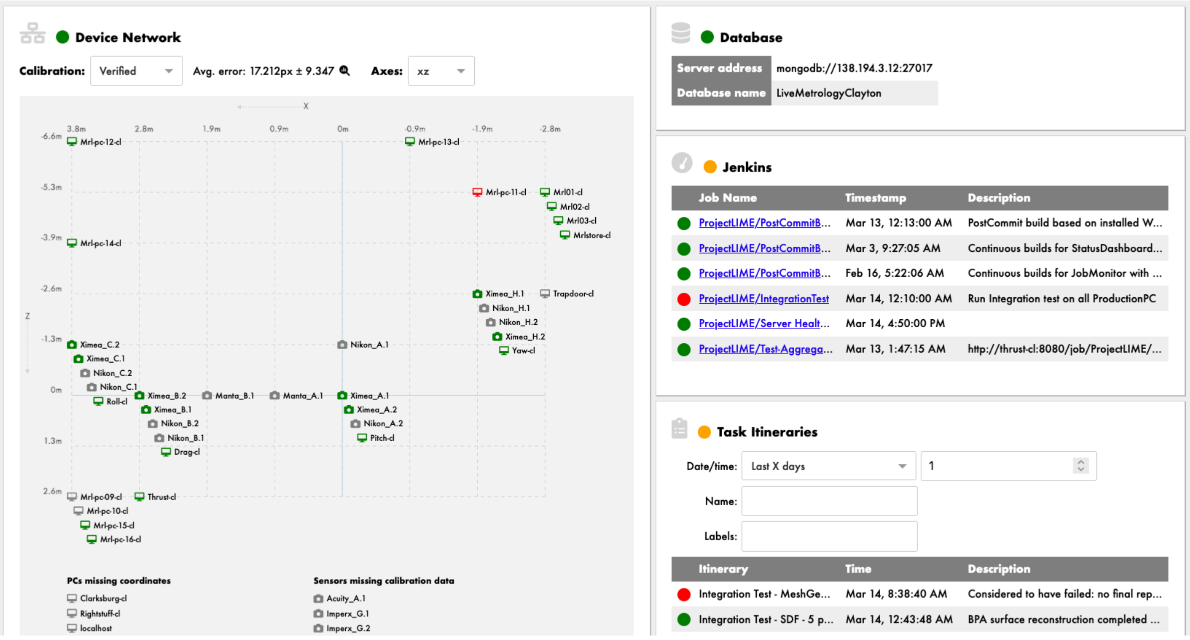

Figure 4. The dashboard application (left: device network status, right-top: database status, right-middle: Jenkins job status, bottom-right: itinerary status)

Device icons can be hovered over to obtain more information (Figure 5), including, for sensors, their calibrated coordinates, calibration error statistics, and screenshots of their last capture.

We can quickly ascertain the status of the database (Figure 4, right top) continuous integration system (Figure 4, right middle) and recently completed task itineraries (Figure 4, right bottom). Hovering over items in the itinerary panel allows us to see the full history of any reports related to this itinerary, which is important as itineraries contain many interdependent tasks that are often asynchronously distributed across the processing network.

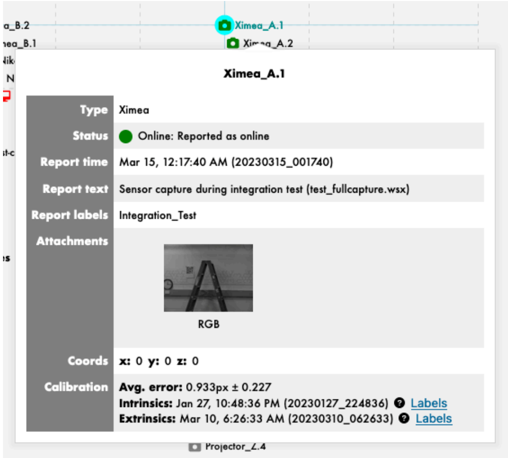

Figure 5. Detailed information about a device, shown on mouse hover

Several attachment types are supported for status reports, including Uniform Resource Locator (URL)s; images, such initial sensor captures or rendered output of surface reconstructions; MetaCloud IDs (see section 3.2); or JavaScript Object Notation (JSON) documents. Images can be submitted to the API as either Base64 encoded Portable Network Graphics (PNG) data or uploaded directly to a desired storage location and referenced as a URL. Images uploaded as Base64 are converted to PNG files and stored in a configurable location in the application's settings.

The status dashboard is built as a standard modern web application, with a back-end REST API built on the Sanic Python framework (Sanic Community Organisation, 2018) and MongoDB (MongoDB inc., 2023) for data storage, and a front-end built using the React Javascript framework (Meta, 2023). It should be noted that the use of MongoDB is a constraint imposed by a key project stakeholder, and that the team has plans to move away from this due to the prohibitive nature of the license (SSPL). A Workspace plugin has been built to enable straightforward submission of status reports from Workspace workflows for standard data types, making it trivial to add status reporting to existing capabilities in the MRL.

3.2. Point-cloud provenance

To help address challenges diagnosing issues in large in point cloud result datasets, and reproducing them in the case of data loss, the team implemented a sophisticated point cloud provenance framework termed "MetaCloud". Point clouds generated by the data acquisition network are aggregated into a single MetaCloud and associated with meta data. MetaClouds flow downstream through the processing network, treated as either a large "point cloud soup", or as a "cloud of clouds", depending on the requirements of the process in question. Each time a MetaCloud is processed, a new MetaCloud is generated, with a reference back to its parent MetaCloud(s) as well as to a record of the process. This allows us to keep track of the inputs used to generate each cloud, as well as the processing history, which can form complex directed acyclic graphs.

For example (Figure 6), in the case of visual quality inspection, initial point clouds can be generated by executing a CSIRO SDF 3D reconstruction process for each pair of sensors in the MRL. We could aggregate the entire set into a single MetaCloud. This MetaCloud could then be denoised, either being treated as one large point cloud, or denoised component-wise. In either case, a new MetaCloud is created, composed of the denoised point cloud(s). This denoised MetaCloud could then be tiled and further reduced for interactive visualisation.

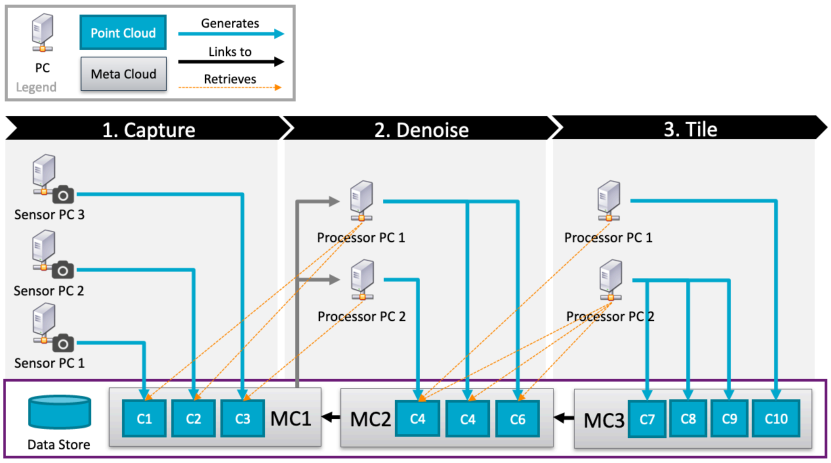

Figure 6. Example of how MetaClouds can be constructed and processed using the distributed compute network

For MetaCloud storage, we have devised a modular data store that indexes point data on a network drive in a custom Hierarchical Data Format version 5 (HDF5) format that facilitates retrieval of sub-regions of point clouds (PLY is also supported) and records the associated meta data/processing history in a database. Both the point storage and metadata storage components can be swapped with alternatives to meet site-specific storage requirements. MetaClouds can be retrieved from the store over using the *GetMetaCloud* operation in the MetaCloud Workspace plugin. Alongside the ability to read/write and interrogate MetaClouds, the MetaCloud Workspace plugin includes a sophisticated Graphical User Interface (GUI) application, *MetaCloud Inspector*, which allows operators to discover, inspect and visualise MetaClouds (Figure 7).

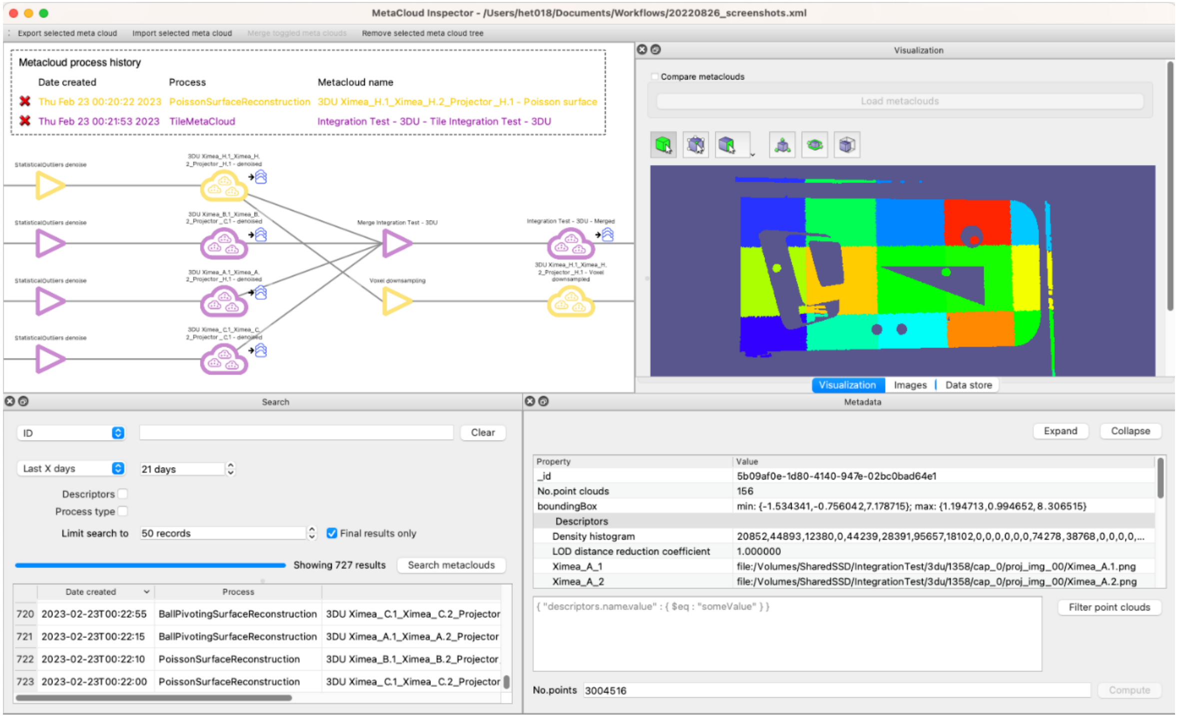

Figure 7. MetaCloud Inspector application (top-left: the MetaCloud dependency graph, top-right: visualization of currently selected MetaCloud, bottom-left: search panel for identifying MetaClouds, bottom-right: meta data panel)

The MetaCloud Inspector application allows operators to quickly locate a dataset of interest using the search panel (Figure 7, bottom-left), then use the interactive process history graph (Figure 7, top-left) to select a MetaCloud for interactive visualization (Figure 7, top-right) and meta-data inspection (Figure 7, bottom right). Search parameters include time range, process type, ID or label, and an option is provided to include "Final Results Only" – those that are not inputs downstream. Multiple results can be viewed simultaneously, allowing comparisons to be made not only between stages within a dataset, but across datasets.

In the process history view, the triangle graph nodes represent "Processes" that have generated one or more MetaClouds, while the "cloud of clouds" icons represent the MetaClouds. Clicking on an icon in this view will show the inputs/descriptors in the Metadata panel (Figure 7, bottom right), including shortcut links to related MetaClouds. Image descriptors are shown as links that, when clicked, display in the "Images" tab.

Results can be visualized using different colouring schemes, such as depth, point clouds index (seen in Figure 7, top-right), Red-Green-Blue (RGB) state, or by a custom metadata descriptor (e.g., all component point clouds with a 'curvature' metadata value > 5 should be colored blue).

4. DISCUSSION AND CONCLUSION

Applying the technology underlying the Mixed Reality Lab to multiple research projects in multiple different environments has allowed our team to explore its effectiveness as a unified platform for the development of digital twin systems involving visual inspection. Implementing digital twin solutions using the platform has led to the discovery of multiple higher order challenges relating to the digital twin's awareness of its own components and capabilities, solutions to which are paramount if any implementation is to function appropriately in a real-world setting without the ongoing support of experienced engineering resources. The

MetaCloud system and Status Dashboard solutions presented in this paper represent flexible and extensible approaches to solving these challenges that appropriately consider the limited engineering resources available to the organisations looking to adopt digital twin technologies. These utilities help make the MRL technology suite an attractive offering for research teams or organisations looking to implement digital twin systems.

REFERENCES

Cao, Z., Hidalgo, G., Simon, T., Wei, S.-E., Sheikh, Y., 2018. OpenPose realtime multi-person 2D pose estimation using Part Affinity Fields, in: ArXiv Preprint ArXiv1812.08008.

Cleary, Paul W., Thomas, D., Hetherton, L., Bolger, M., Hilton, J. E., Watkins, D., 2020. Workspace: A workflow platform for supporting development and deployment of modelling and simulation, Mathematics and Computers in Simulation, 175, p25-61, ISSN 0378-4754.

Culjak, I., Abram, D., Pribanic, T., Dzapo, H., Cifrek, M., 2012. A brief introduction to OpenCV, 2012 Proceedings of the 35th International Convention MIPRO, ISBN 978-953-233-068-7.

Czichos, H., Saito, T. and Smith, L. E., 2011, Springer Handbook of Metrology and Testing. p10. ISBN 978-3-642-16640-2.

Endsley, M. R., 1995, Toward a theory of situation awareness in dynamic systems. Human Factors 37 (1), 32–64. doi:10.1518/001872095779049543

Harrison, Simon M., Cohen, Raymond C. Z., 2021. Mechanical Digital Twinning of the human body in the workplace for reduced injury risk and improved health. 24th International Congress on Modelling and Simulation (MODSIM) 2021; p183-189.

Hetherton, L., Bolger, M., Zhang, Z., Oakes, N., Watkins, D., Cleary, P. W., 2019. Workspace for Industry 4.0. 23rd International Congress on Modelling and Simulation (MODSIM) 2019, p449-455.

Jenkins, 2023, <https://www.jenkins.io/> Accessed 01.03.2023.

Meta, 2023, React <https://reactjs.org/> Accessed 01.03.2023.

MongoDB Inc, 2023, MongoDB <https://www.mongodb.com/> Accessed 01.03.2023.

Rebo, M., Brandner, M., 2005, International Workshop on Robotic Sensors: Robotic and Sensor Environments, 2005, doi:10.1109/ROSE.2005.1588327.

Safe Work Australia, 2022, Safer, healthier, wealthier: The economic value of reducing work-related injuries and illnesses, Accessed 01.03.2023.

Sanic Community Organisation 2023, Sanic <https://sanic.dev/en/> Accessed 01.03.2023.

Stainlay, M. L., Poropat, G. V., 2005. A method for automated 3d imaging. WO2006024091A1 <https://patents.google.com/patent/WO2006024091A8/3Den>.

Watkins, D., Bolger, M., Hetherton, L. Zhang, Z., Xie, P., Harrison, S. M., Cohen, R. C. Z., 2021, Using Workspace to implement Digital Twins in the Mixed Reality Lab, 24th International Congress on Modelling and Simulation (MODSIM) 2021, p191-196.

Wu, Y., Kirillov, A., Massa, F., Lo, W.-Y., Girshick, R., 2019. Detectron2.