Artificial Intelligence in Makerspaces – Repurposing industry applications to serve makerspace needs

6th International Symposium on Academic Makerspaces

ISAM
2022
Paper No.:
84

Domenic DiCarlo[1], Nicholas Greenfield[2], and Amit S. Jariwala[3]

[1]Domenic DiCarlo; School of Mechanical Engr., Georgia Institute of Technology; e-mail: ddicarlo3@gatech.edu
[2]Nicholas Greenfield; School of Mechanical Engr., Georgia Institute of Technology; e-mail: ngreenfield@gatech.edu
[3]Amit S. Jariwala; School of Mechanical Engr., Georgia Institute of Technology; e-mail: amit.jariwala@gatech.edu

Introduction

Artificial Intelligence, in modern understanding, is the use of computers and computing power to perform tasks previously known to be only possible by a human. Until recently, there has not been enough processing capability to allow for these activities. Now, developments in computing capabilities have allowed for the prospering of artificial intelligence (AI). AI has had a pronounced effect on the manufacturing industry by allowing it to make sense of the large amounts of data generated by sensors and other monitoring devices. This data, known as Big Data, is too dense to be processed by solely human intuition. However, by using machine learning (ML), a subset of AI, trends, and patterns can be discovered and leveraged to the industry's benefit. Any manufacturing setting puts its occupants in a certain degree of danger. Mitigating that danger is critical in any scenario, but especially in the environment of a makerspace. Makerspaces, by nature, are open to a wider spectrum of users with diverse skill sets and backgrounds than a typical industrial environment. Ensuring that even the most novice user can safely access, train, and utilize the equipment in the space is paramount. One of the many objectives that makerspaces serve is to lower the barriers of entry and access to prototyping/fabrication equipment. This paper presents novel approaches by which makerspaces could leverage the wide accessibility of AI tools, IoT sensors, and computation power to better achieve their mission and serve their stakeholders.

Background

Over the last decade, the popularity of makerspaces has increased dramatically. Various factors, such as the advent of rapid prototyping and changes in academic culture, have allowed for their prosperity. Such growth introduces challenges that may have been unimportant otherwise. For example, makerspaces see many people of many different backgrounds [1]. The diversity of these backgrounds, while advantageous in bringing forth different ideas and solutions, creates issues when establishing a baseline body of knowledge shared by the entire space. Without knowing how much everyone knows, it can be difficult to catch cases that fall through the cracks. In the best of cases, this can lead to broken equipment and reduced functionality to allow for repairs; in the worst cases, serious injury can be incurred.

The manufacturing industry has largely the same problem to deal with. Different workers from different walks of life all need to be able to use machinery safely and effectively. This is typically done through comprehensive training and careful education, both formal and practical, on best practices. Makerspaces do this as well; however, it is more difficult for them to reach the same level of comprehension due to the variety of backgrounds they encounter. While in manufacturing, the same part or component, or mechanism is made repeatedly, makerspaces see a much wider range of use. The best practice for one project may not necessarily be the best practice for another. This makes it more difficult to guide users on how the equipment should be used in the most suitable way.

One means of expanding a makerspace's body of knowledge is through the use of big data and machine learning. By collecting information from machines and using that information to train models to recognize inefficient or damaging behavior, many of the issues preventing more generalized use can be addressed. The manufacturing industry has made ample use of machine learning in various applications, including condition monitoring and predictive maintenance, image detection, and energy efficiency. With minor alterations, makerspaces could see a similar benefit from the technology.

There are several provisions in place to ensure the safety of users. Advanced fabrication equipment like laser cutters, CNC machines, etc. includes safety interlocks and emergency stops. However, these safety features are not standard on all equipment used in makerspaces. Rigorous training sessions could prepare the end-user for safe use of the machine but cannot guarantee zero accidents. Makerspace staff also provide a line of defense against safety incidents by serving to monitor operations and tool usage within the space and correcting the user when necessary. However, humans are prone to making errors and hence a non-intrusive support system, comprised of sensors and paired with machine learning models can be very helpful in a makerspace.

In addition to assisting makerspaces in providing a safe environment for users, artificial intelligence can help makerspaces operate with greater efficiency. Machines require regular maintenance to perform their duties without incident. Without knowing the exact state of the equipment, maintenance must be done based on when the previous maintenance was done and the aggregate operational time of the device. While this ensures the machine stays operational, it may result in servicing being done more often than strictly necessary. This double affects the makerspace, both by consuming resources more frequently than required and requiring the machine to be shut down more often than needed. Machine learning enabled predictive analytics can inform staffers about the state of the machine. In doing so, it allows them to make more informed decisions about when to perform maintenance. This saves resources involved in the maintenance and keeps the functionality of the space as high as possible.

The authors leveraged the Vertically Integrated Projects (VIP) program at Georgia Tech to develop intelligent systems and tools to manage and maintain makerspaces. Undergraduate students from different majors and academic standing receive an opportunity to work with graduate students and faculty while receiving elective research credits. This paper presents some of the work conducted by these teams in the past few semesters. More information about this team is available here: https://www.vip.gatech.edu/teams/vx4

Literature Review

An important aspect of any manufacturing industry is maintaining the machinery used to create a product. Maintenance takes time and money away from production, so minimizing the amount of maintenance needing to be done is in the interest of all manufacturing industries. One method that allows manufacturers to better determine how and when maintenance needs doing is condition monitoring. Condition monitoring involves using sensors to actively track the operation of a machine and assess whether faults have taken place or components are becoming worn. In their doctoral thesis [2], Mabrouka Baqqar utilized data collected from operating parameters to detect broken tooth faults in gearboxes. To do this, three different artificial intelligence models were trained to recognize deviations in current, load, and temperature from normal. These AI models were a general regression neural network, a back propagation neural network, and an adaptive neuro-fuzzy inference system. Each was successful in their training and recognition of abnormal behavior.

Gearboxes are necessary for many industrial applications to suit the needs of the process. Due to their ubiquity, it is important to know when faults appear to address them as quickly as possible to minimize machine downtime. Most applications use vibratory or acoustic data to detect these faults, requiring additional sensors [3][4][5]. Consequently, additional complexity and cost are introduced to the machine. Baqqar's approach is unique in this respect. By leveraging the operating parameters as opposed to using external sensors, they accomplish the goal of fault detection without additional complication.

While makerspaces don't bear the burden of per-unit production cost, they do share the concern for minimizing machine downtime. The longer a machine is down, the longer the space is not operating at full functionality. For this reason, makerspaces have a vested interest in this kind of technology. Furthermore, an approach that does not require modification of equipment or the purchasing of additional sensors is well suited for a makerspace environment, as it minimizes the overhead required. This allows those costs to be utilized elsewhere, such as by increasing capacity or hiring additional staff. Additionally, by detecting faults as they occur, higher consequence failures are avoided, thereby improving the safety of the space.

Another area of manufacturing that has benefitted from the use of machine learning has been quality control. Defects of some components can be microscopic while still having an enormous impact on the final product, and technicians can miss such small deviations. To combat these errors, the advent of image detection has been a boon. This technology has been applied by Weimer et al. [6]. In their paper, the team utilized a convolution neural network (CNN) for industrial inspection for defects, tweaking the setup along with the width and depth parameters. Their choice to use a CNN here was not arbitrary. Previous applications of image detection in this field require manual feature selection. For known defects, this works fine; however, if new kinds of defects appear or are present in various orientations or sizes, a new feature set must be defined. In these events, a CNN has the advantage of not requiring a manual definition of features. Instead, the CNN "automatically generates meaningful features…with minimal human interaction."

Unlike the manufacturing industry, makerspaces have little specific concern with defects in parts. Due to the variety of projects that makerspaces witness, it is uncommon for microscopic flaws in a part to have an effect on the project being worked on. However, makerspaces can still benefit from technologies like image detection for other applications. Training of staff typically requires a demonstration of proper use of a piece of machinery, with proof given via some part made to certain specifications. In these instances, image detection could be able to recognize when parts are out of spec or take note of signs that improper methods have been used. Alternatively, image detection can be used to identify improper machine use in real-time and alert staff to address the issue. Practices like these would serve a makerspace by mitigating the risk of broken machinery or harmed staff or users.

To this last point above, machine learning is used in some manufacturing environments to sense the surroundings of industrial robots. Industrial robots typically operate on strict paths with little ability to alter their movements to avoid obstacles. The addition of sensory capabilities would allow them to work better in collaboration with humans and prevent workplace accidents. This sensory ability is referred to as machine vision by Golnabi and Asadpour [7]. Furthermore, Kuts et. al and Bexten et. al. compare and evaluate different means of detection and planning related to machine vision for industrial robots in their works [8] [9]. The parallels between machine vision applied to industrial robots and its applications for makerspaces are clear. Makerspaces regularly involve people interacting with potentially dangerous equipment. To be able to sense when people are in harm's way and dissuade users from continuing the unsafe operations in such scenarios would improve the safety of these environments.

AI has seen commercial success in industrial applications outside of the manufacturing sector. Hassanaly discusses various inclusions of AI/ML in the construction industry, as well as the benefits it provides for its use as well as the potential risks and how they can be mitigated [11]. Similarly, Voxel AI, a tool by which safety incidents can be mitigated and reported, saw success in its implementation in Americold's California distribution center [12]. These applications function by observing work through a system's security cameras and identifying unsafe behavior, such as the improper lifting of pallets or speeding equipment.

The ability to process large amounts of data generated from equipment and find patterns that would otherwise go unnoticed allows machine learning to assist people outside of the traditional manufacturing industry. In the domestic area, Sense technologies allow homeowners to monitor their electricity usage and make more informed decisions about how power is being consumed [10]. Sense integrates into the circuit breaker of a building to monitor energy use. Through machine learning, it gains the ability to determine which energy signatures are related to household devices. Furthermore, it can recognize patterns in usage and alert the owner to abnormal appliance behavior. In doing so, it can save the homeowner money by alerting them to issues before more disastrous circumstances present themselves.

Though not typically designed for it, technologies like Sense could serve the makerspace community. Makerspaces house many pieces of equipment that require relatively large amounts of power to operate. Being able to track which devices are used frequently can enable a space to make more informed decisions about what equipment is popular, allowing for more informed decisions regarding the expansion of functionality. Additionally, as with [2], detection of uncharacteristic behavior indicative of faults would be beneficial. Moreover, having a better sense of how power is being consumed gives makerspaces the ability to be more energy conscious in their purchasing. In doing so, money is saved, which could be reallocated to other services as needed.

The existing applications of AI/ML from literature are all designed for applications within the industry or commercial settings. Previous researchers [13-16] presented a thorough literature review and inspiring examples of how IoT sensors could be used in makerspaces to collect data. However, the current knowledge gap is how to extend the applications of sensors to collect data and leverage AI/ML for useful purposes to better serve the makerspace needs. The subsequent sections present two case studies and future work that could help makerspace managers/administrators, or developers leverage these novel technologies to support makerspaces.

Case Study One: Detection of Improper Bandsaw Use

Makerspaces often see many users in a single day. While ample utilization of the space can hardly be complained about, it often leads to issues when the number of users outpaces the supervisors' ability to effectively manage safety. To address this issue, a team pursued the goal of improving makerspace safety by detecting when a user's hands were dangerously close to the blade of a bandsaw without human assistance.

To capture images, the team used a cell phone recording video positioned above the bed of the bandsaw pointed downwards. The video would then undergo preprocessing, involving grayscaling the images, reducing the resolution, and increasing the contrast. Each frame of the video would then be subject to a machine learning model, and safety judgment would be assigned. In the event of unsafe usage, the model would alert an on-staff Prototyping Instructor (PI) to the transgression. After evaluating several options, the team decided to use a Microsoft HoloLens to inform the PI. A flowchart illustrating the process is shown in **Fig. 1.**

Initially, the team attempted to train a CNN by taking videos of the bandsaw being used both properly and improperly. These videos were labeled as either "safe" or "unsafe" and given to the model to train it. However, instead of learning what aspects of the footage constituted safe versus unsafe behavior, the CNN only memorized the training set and made safety judgments based on the sequential position of frames. When tested in a live setting, the CNN failed to generalize and incorrectly identified all use of the bandsaw as safe. From there, the team changed its approach. It implemented the use of YOLO, an "extremely fast and accurate" object detection system [17], using it to identify the position of the user's hands and the location of the bandsaw blade.

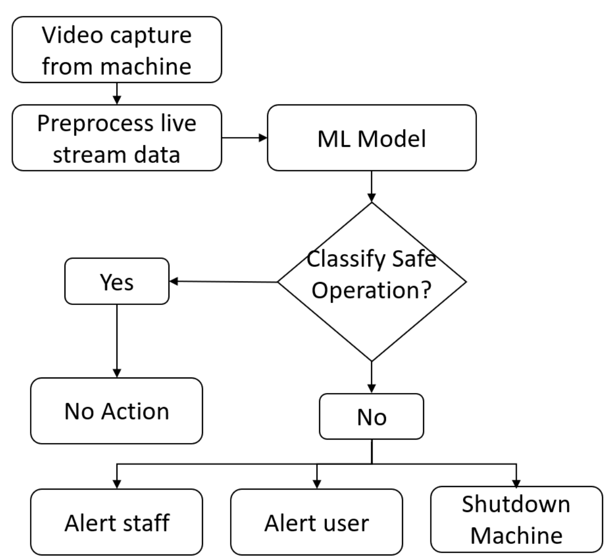

Fig. 1: Flowchart describing the information flow

YOLO was trained to recognize these items by annotating video footage using CVAT, as shown in **Fig. 2**. The annotated images were then uploaded to Roboflow to generate a dataset to be used to train the YOLO model. The dataset was uploaded to Google Colab to attempt to train the YOLO model.

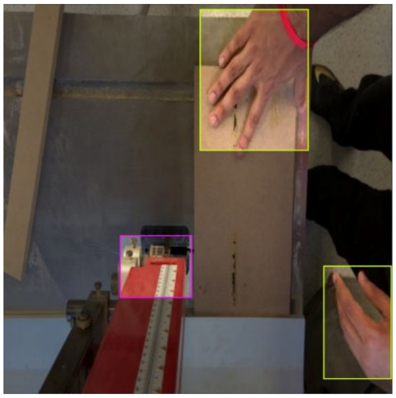

Fig. 2: Annotations are drawn using CVAT.

This project is currently a work in progress. Next steps would involve implementing the model in a real-time environment, with logic in place to evaluate whether the relative position of the blade and hands is unsafe. The research team engaged with the Institute's Environment & Health Safety office as well as the Institute Risk Management leadership to explore the feasibility of implementing the proposed approach to augment safety within the Institute's makerspace. Both the groups expressed strong support for this approach as it enhances user safety by improving the duty of care provided by the makerspace to its users. However, they recommended some critical design aspects that should be central to the final implementation of this solution. These were:

1. Avoid physical alterations to the equipment: Avoid impairing the regular operation of the machine and ensure proper risk assessment is conducted to make sure any sensors/actuators/alarms added on or around the equipment do not cause any harm to the user in case of unintended failures of the system.
2. Be mindful of user privacy and only collect/analyze data that is essential
3. Ensure that the user/staff notification system does not distract the user or other users in the space.
4. Instead of a binary gauge of detecting safe vs unsafe operations, identify ways to track near misses which would help train future users
5. Implement ways to educate the end user that the AI support elements are only tertiary lines of defense against accidents and users are solely responsible to ensure their own safety.

The primary advantage of the proposed machine vision-based safety monitoring system is that it can be placed outside of the bounds of the machine and placed in such a way that user's face is not recorded. If the technology were to be integrated into the machine, it must be ensured that the safe use of the machine is not otherwise compromised. In their paper, Anastasi et. al. discusses how various standards for safe design and use of machinery from the European Union apply to AI/ML [18]. Furthermore, they examine how these standards might be revised to better cover machinery with embedded artificial intelligence. While such standards are for the EU market, they provide guidelines for the general implementation of AI/ML in the manufacturing industry and would be relevant to the research discussed in this case study.

Addressing the third point above, research into design of safety alerts/ alarm has been done; in their paper, Schlesinger et. al. investigate how alarm volume affects the cognition of clinicians performing patient-related duties while subject to distractors [19]. Similarly, Rayo et. al. studied the use of timbre to aid nurses in recognizing the category and urgency of alarms [20]. While this research is focused on the medical field, it is still applicable to users of heavy machinery; both groups require the participant to perform actions, sometimes more than one simultaneously, with care, the penalty for which could be bodily harm to themselves or others.

In addition to designing alarms that don't create unnecessary distraction, the concept of implementing tiered zones of safety is being discussed. Such zones would allow for failure modes that don't involve harm befalling the user. Inclusion of such zones would be paired with lower-level alerts so that the danger of the use is accurately communicated to the user. This would serve to address the fourth point of the list above.

While this research is focused on the use of a bandsaw, the applications of it can be generalized to other pieces of equipment. Any machinery that involves the user interacting with or operating close to a cutting part would benefit from the research done here.

Case Study Two: Characterization of Cut Path Aggressiveness

CNC mills are often intimidating machines for new users in makerspaces. Adequate training itself is no substitute for hands-on experiences with trial and error and hence it is possible for the machine to be used improperly, or rather in a non-optimal fashion. In such events, the best-case scenario involves tool wear and inefficient work, resulting in lost time or broken equipment; in the worst cases, the health and safety of the operator are at stake. For a makerspace, where the spectrum of the base knowledge and skills of the user varies widely, the challenge presented by the CNC can drive away potential users. In order to improve the safety of the user and improve the accessibility of the makerspace, an ML model was developed to leverage onboard sensors to classify how aggressively a CNC is being used, i.e., how intensely the cutting bit is being engaged.

A flow chart describing the actions taken is shown in **Fig. 3.** Aggressiveness was assessed using ranges of feed-rates taken from G-Wizard, an online tool for choosing machining parameters. To gather training data, the author ran cuts of simple geometry, recording spindle power data as well as loads in the x, y, and z directions. Additionally, individual loads were combined to find the total load magnitude. This data was processed to remove unnecessary portions, such as startup and shutdown. Frequency features were generated through a Continuous Wavelet Transform (CWT), then reduced to increase the speed of the model. Reduction was done using Pearson Correlation and Mutual Information (MI) scores. After selection, the remaining features were used to train several different ML models. These models include Logistic Regression, K-Nearest Neighbors (KNN), Decision Trees, Random Forests, and Multi-Layer Perception (MLP) Neural Network. Because each of these models uses different hyperparameters, optimization was done by comparing F1 scores after changing various values.

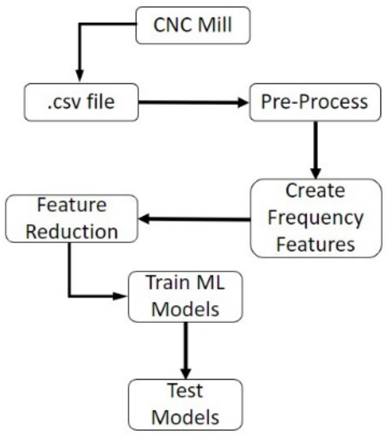

Fig. 3: Flow chart describing the actions taken by the author in conducting the research into CNC aggressiveness.

Once the models were trained, they were subject to two tests. The first test only involved different cutting geometry, while the second also varied the cutting parameters like width and depth of cut. This second test was done to mimic the different conditions a CNC machine would be subject to in a makerspace environment, where cutting parameters change depending on the project. In each test, the models were compared using the area under the Receiver Operator Characteristic (ROC) curves. Under the conditions of a makerspace, the models performed poorly, so a Principal Component Analysis (PCA) was performed to reduce the noise. Using a PCA with n=2, Logistic Regression functioned the best. **Table 1**, shown below, shows how including this PCA improved the models' predicting abilities.

Table 1: Areas under ROC curves for different ML models at different cut aggressiveness, before and after applying PCA. The cells in green mark the instances of PCA improving the prediction.

		Logistic Regression	KNN	Decision Tree	Random Forest	MLP
No PCA	Conservative	0.404	0.538	0.454	0.538	0.506
	Optimal	0.389	0.559	0.474	0.358	0.540
	Aggressive	0.588	0.621	0.602	0.704	0.437
PCA n=2	Conservative	0.788	0.623	0.601	0.637	0.478
	Optimal	0.573	0.479	0.490	0.500	0.497
	Aggressive	0.785	0.600	0.640	0.665	0.566

The results show that it is possible to predict cut aggressiveness using just the onboard sensors of a CNC. Makerspaces have ample use for this technology in order to reduce the barriers of entry to the machine as well as improve the overall safety of the space, as is shown in **Fig. 4**. A more varied training set is needed to improve the accuracy of the model, including tool wear, workpiece material, and a wider spectrum of feed rate and width of cut.

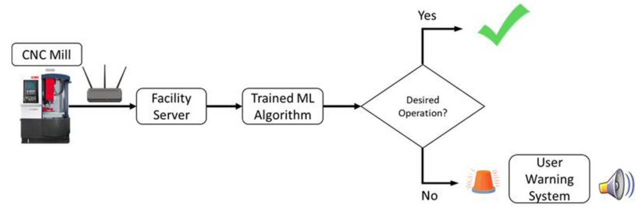

Fig. 4: Flow chart depicting the potential use cases of the aggressiveness classification.

While this specific research focuses on identifying aggressiveness from machine parameters, deeper implications exist. For example, detecting the work done by the machine through the built-in sensors could be used for the identification of materials. From there, recommendations could be made for cutting parameters, or warnings could be given for improper or dangerous use before dire consequences are encountered. Similarly, with a given material, the improper setup could be noted before danger is encountered. This would serve to both improve the safety of the space as well as introduce newer users to higher-level concepts of CNC milling as they encounter them. By having a specific event with which to remember them, users would be more likely to remember the concepts introduced than if they were abstractly mentioned during regular training.

Conclusion

Makerspaces are becoming commonly adopted across educational institutions as hubs for cross-disciplinary collaborative learning and innovation. Due to the cross-disciplinary nature, users with diverse backgrounds and skills visit the makerspace. In order to make makerspaces more accessible to this vast audience, it is necessary to lower the barriers to entry while enhancing the overall safety of the space. Recent developments in computing capabilities have made democratized access to AI tools and this paper presented specific case studies where it makes sense to leverage AI as a tool to support makerspace operations. Existing literature suggests the growing use of AI to support manufacturing operations for industrial and commercial needs. This paper presented inspiring examples of how free-to-use AI tools could be used to develop technologies specific to makerspaces. An approach on how to use ML to confirm the user's intent in machining and provide real-time actionable insight to the end-user was also presented. This approach can be used to make the machining operation safer for the user thus increasing its accessibility.

Acknowledgment: This material is based upon work supported by the Department of Energy / National Nuclear Security Administration under Award Number(s) DE-NA0003921.

rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of the authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

References

[1] Bouwma-Gearhart, J., Choi, Y. H., Lenhart, C. A., Villanueva, I., Nadelson, L. S., & Soto, E. (2021). Undergraduate students becoming engineers: The affordances of university-based makerspaces. Sustainability, 13(4), 1670.

[2] M. Baqqar, "Machine performance and condition monitoring using motor operating parameters through artificial intelligence techniques," Ph.D dissertation, School Of Comp. and Eng., Univ. of Huddersfield, Huddersfield, 2015.

[3] Bohara, N., & Sakthivel, G. (2017). Carbide coated insert health monitoring using machine learning approach through vibration analysis. International Journal of Prognostics and Health Management, 8(2).

[4] Umbrajkaar, A. M., Krishnamoorthy, A., & Dhumale, R. B. (2020). Vibration analysis of shaft misalignment using machine learning approach under variable load conditions. Shock and Vibration, 2020.

[5] Balachandar, K., Jegadeeshwaran, R., & Gandhikumar, D. (2020). Condition monitoring of FSW tool using vibration analysis–A machine learning approach. Materials Today: Proceedings, 27, 2970-2974.

[6] Weimer, D., Scholz-Reiter, B., & Shpitalni, M. (2016). Design of deep convolutional neural network architectures for automated feature extraction in industrial inspection. CIRP Annals, 65(1), 417-420.

[7] Golnabi, H, Asadpour, A. (2007). Design and application of industrial machine vision systems. Robotics and Computer-Integrated Manufacturing, 23(6), 630-637.

[8] Kuts, V, Otto, T, Tähemaa, T, Bukhari, K, & Pataraia, T. "Adaptive Industrial Robots Using Machine Vision." *Proceedings of the ASME 2018 International Mechanical Engineering Congress and Exposition. Volume 2: Advanced Manufacturing*. Pittsburgh, Pennsylvania, USA. November 9–15, 2018. V002T02A093. ASME. https://doi.org/10.1115/IMECE2018-86720

[9] S. Bexten, J. Saenz, C. Walter, J. -B. Scholle and N. Elkmann, "Discussion of using Machine Learning for Safety Purposes in Human Detection," *2020 25th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*, 2020, pp. 1587-1593, doi: 10.1109/ETFA46521.2020.9212028.

[10] C. McGarry, "*How Does Sense Detect My Devices?*" https://blog.sense.com/articles/how-does-sense-detect-my-devices/ (accessed June 15, 2022).

[11] Hassanaly, D., *"Artificial Intelligence and Sensor-based Health and Safety Monitoring"* https://www.gihub.org/resources/showcase-projects/artificial-intelligence-and-sensor-based-health-and-safety-monitoring/ (accessed July 14, 2022).

[12] *"Transforming Safety at Food Distribution Warehouses: Reducing Accidents and Operational Costs"* https://www.voxelai.com/case-study/#top (accessed July 14, 2022).

[13] Banerjee, R., Ferron, L. M., Bailey, D., Imam, R., & Jariwala, A. S. IoT in Makerspaces-Review of Usage and Future Potential 4th International Symposium on Academic Makerspaces.

[14] Cooke, M. N., & Charnas, I. C. (2016). The Value of Data, Metrics, and Impact for Higher Education Makerspaces. In Proceedings of the 1st International Symposium on Academic Makerspaces (No. 37, p. 180).

[15] Imam, R., Patel, P. B., Ferron, L., & Jariwala, A. S. Development and Impact of a Data Collection System for Academic Makerspaces 2nd International Symposium on Academic Makerspaces.

[16] Sung, G., Yoo, T., Chng, E., Yang, S., & Schneider, B. (2021). Automated Tracking of Student Activities in a Makerspace Using Motion Sensor Data. In Proceedings of the 14th International Conference on Computer-Supported Collaborative Learning-CSCL 2021. International Society of the Learning Sciences.

[17] *"YOLO: Real-Time Object Detection,"* https://pjreddie.com/darknet/yolo/ (accessed June 7, 2022).

[18] Anastasia, S., Madonna, M., Monica, L. (2021). "Implications of embedded artificial intelligence – machine learning on safety of machinery," *Procedia Computer Science*, 180, 338-343.

[19] Schlesinger, J., Miller, S., Nash, K., Bruce, M., Ashmead, D., Shotwell, M., Edworthy, J., Wallace, M., Weinger, M. "Acoustic features of auditory medical alarms—An experimental study of alarm volume", The Journal of the Acoustical Society of America 143, 3688-3697 (2018) https://doi.org/10.1121/1.5043396

[20] Michael F. Rayo, Emily S. Patterson, Mahmoud Abdel-Rasoul & Susan D. Moffatt-Bruce (2019) Using timbre to improve performance of larger auditory alarm sets, Ergonomics, 62:12, 1617-1629, DOI: 10.1080/00140139.2019.1676473