

Short Communication

Estimating household water storage from images: A machine learning approach

Chad Staddon^{a,*}, Ajmal Shahbaz^b, Syed U. Yunas^b, Lyndon Smith^b, Geraint Burrows^c, Sayed Mohammed Nazim Uddin^{d,e} and Lucy Whitley^c

^a Centre for Water, Communities & Resilience, University of the West of England, Bristol BS16 1QY, UK

^b Centre for Machine Vision, Bristol Robotics Lab, University of the West of England, Coldharbour Lane, Bristol, BS16 1QY, UK

^c Groundwater Relief, Dartington Hall, Totnes, Devon TQ9 6EQ, UK

^d IMDEA Water Institute, Avda. Punto Com 2, 28805 Alcalá de Henares, Madrid, Spain

^e Center for Climate Change and Environmental Health (3CEH), Asian University for Women, 4000, Chattogram, Bangladesh

*Corresponding author. E-mail: chad.staddon@uwe.ac.uk

 CS, 0000-0002-2063-8525

ABSTRACT

Whilst much research focusses on challenges related to achieving SDG 6.1 (universal and equitable access to safe and affordable drinking water), there has been less attention to challenges of safe transport, storage and use of collected water. In particular, there are relatively few high-quality datasets quantifying the number and volume of water containers used by households for such purposes. This paper reports results from the application of machine learning (ML) techniques to a database of images of domestic water storage collected during 2022 as part of an initiative to improve water supply in southern Bangladesh. Because the number of different water container types was relatively small, it was possible to train an ML algorithm to identify water containers and estimate water storage with greater than 90% accuracy. These results have allowed the rapid creation of a unique high-quality, high-resolution dataset describing water storage quantitatively in a study community. This dataset includes data quantifying the number of vessels as well as their individual and aggregated water storage volumes. The paper discusses policy implications for the study location specifically before concluding with suggestions for the inclusion of this sort of analysis in ongoing studies of household and community scale water insecurity.

Key words: Bangladesh, household water storage, machine learning, SDG 6.1

HIGHLIGHTS

- Machine learning can speed up the processing of images of water storage taken during conventional household water access questionnaires.
- Accurate quantification of available water storage can assist in water services programme planning.
- Machine learning-assisted analysis of over 800 images collected during household research in Bangladesh in 2022 showed that the average storage per person is only 2.5 L.

INTRODUCTION

In water-insecure communities, the humble (often yellow) ‘jerry can’ is a familiar sight, often alongside a jumble of other (usually plastic) water storage containers (Figure 1). In 2022, for example, there were over 1.5 billion people worldwide who were dependent on water provisioning arrangements that required collection and transport from remote locations and storage in the home prior to use (GBD 2016 Diarrhoeal Disease Collaborators 2018; Prüss-Ustün *et al.* 2019; Staddon & Brewis 2024). Fetching water from the source to home with these commonly used containers generates immediate risks from the journey itself as well as risks associated with subsequent management of transported and stored water. In addition to chronic musculoskeletal problems caused by repeatedly carrying water even short distances, transporting water manually also carries with it a higher likelihood of acute injury, particularly bone fractures and muscle and ligament damage (Geere *et al.* 2010). Research has also found that fetched and then stored water is very often contaminated with microbiological and other contaminants within 12 h of collection (Opryszko *et al.* 2013; Meierhofer *et al.* 2019; Nurjana *et al.* 2023). Although

This is an Open Access article distributed under the terms of the Creative Commons Attribution Licence (CC BY 4.0), which permits copying, adaptation and redistribution, provided the original work is properly cited (<http://creativecommons.org/licenses/by/4.0/>).



Figure 1 | Variation in size and shape of commonly used domestic water containers (Source: Staddon personal collection).

water containers are a very important aspect of domestic water security around the world, they have received relatively little attention in the research and policy literature (Wani *et al.* 2022; Staddon & Brewis 2024).

Water and Sanitation for Health (WASH) researchers frequently use survey questionnaires to elicit data about type, volume and quality of water storage infrastructure including containers at the household scale. The ability to triangulate survey-derived qualitative data with data generated through machine learning (ML) analysis of collected images offers a number of important advantages. First, there are obvious issues with self-reported data on topics like water storage, including the inability to exactly recall or express storage in numbers of containers and their volumes. Second, because WASH research often involves lengthy questionnaires covering a variety of subtopics, any opportunity to automate data analysis in a subtopic area could help reduce overall survey length and the resulting burden on respondent households. Third, there may be useful information to be generated from comparing self-reported water storage and ML-derived estimates, in terms of a better understanding of cultural and other biases inherent in conventional WASH research. Finally, it may be possible for ML analysis to generate data about related topics of interest including cleanliness of containers, the prevalence of appropriate capping/closing of openings and appropriateness of storage locations – all important factors in water quality management.

The applications of ML-driven object detection are vast but as yet little applied to WASH research. For instance, in autonomous driving, object detection systems play a critical role in identifying pedestrians (Wu *et al.* 2018; Shahbaz & Jo 2021), vehicles (Diwan *et al.* 2023), and traffic signs, ensuring safer navigation (Wang *et al.* 2023a). In the retail sector, they facilitate better inventory management and enable cashier-less checkout experiences (Wang *et al.* 2023b). In healthcare (Wu *et al.* 2018), object detection aids in medical image analysis, from detecting tumours in radiology scans to tracking cells in microscopy images (Chowdhury *et al.* 2023), and in water management, ML-based object detection has been used to track qualitative and quantitative changes in surface water features using satellite imagery (Kashtan & Hnatushenko 2024; Liu & Li 2024).

This paper introduces an artificial intelligence powered algorithm situated at the intersection of computer vision and ML, specifically for assessing the number and volume of water containers from photographs collected as part of conventional WASH research. The primary technical goal here is to utilize automated object detection to correctly identify commonly used water storage containers, such as bottles, buckets, jugs, and more. These data can be used to quickly assess the adequacy of domestic storage volumes in humanitarian and development contexts and to evaluate the impact of specific programme interventions such as the distribution of water containers.

The subsequent sections of the paper are organized as follows:

- Section 2 presents the methodology and discusses key aspects of the ML algorithm used, offering quantitative and qualitative insights into its performance.
- Section 3 presents data on model accuracy in terms of precision and recall statistics.
- Section 4 discusses the results of applying this algorithm to a collection of images of water storage collected in 2022 in Bangladesh.

The final section of the paper discusses some areas for future development, that will be of interest to WASH researchers.

METHODOLOGY USED FOR THE PILOT

In 2022, the UK-based charity Groundwater Relief, the University of the West of England and Asian University for Women in Bangladesh led a study of water infrastructure needs in Teknaf Upazila, located in the southernmost part of Bangladesh. In recent years, this part of Bangladesh has received more than one million Rohingya refugees fleeing government-backed persecution in neighbouring Burma (Myanmar); which has placed great strain on underdeveloped water services infrastructure. This new work was built on earlier studies undertaken by the project team (e.g., Akhter *et al.* 2020; Rafa *et al.* 2020) but specifically sought to capture data about the relative water insecurity experience of host communities and refugee households across wet and dry seasons. During the household survey data collection phase, more than 800 photographs of respondent households' water fetching and storage containers were collected. These images were confirmed to not have identifying data and were made available for the model creation.

The trial algorithm employs YOLO (Jocher *et al.* 2022), or 'You Only Look Once' as an object detection algorithm due to its proven efficiency and efficacy. Imagine if you could look at a photo and immediately point out objects in it, like cars, dogs, or people (water bottles in this study). Now, imagine doing that as fast as flipping through a large stack of pictures. This is essentially what YOLO does; it is a technology that helps computers recognize and find objects in images and videos very quickly and efficiently. YOLO is optimized for fast inference performing object bounding and labelling in a single pass (Fan & Song 2024; Liang *et al.* 2024). Additionally, later versions of YOLO offer a PyTorch implementation with automatic data augmentation, hyperparameter tuning, and pre-trained model support making it highly accessible for both researchers and developers.

Imagine you have an $N \times N$ grid overlay on your photo, like a Noughts and Crosses board but with considerably more squares. YOLO splits each image into a grid, and each square of the grid is responsible for detecting objects that fall within it. This helps the system quickly identify different objects in various parts of the image. In each square of the grid, YOLO predicts two things: the probability of an object being present and the precise location of the object (by drawing a box around it). For example, if there is a water bottle in the photo, YOLO will highlight the water bottle and tell you exactly where it is. YOLO does not just detect objects; it also classifies them. This means it recognizes what each object is – whether it is a dog, a bicycle, or a water bottle. It does this by using a database of known objects (derived during model training) and matching what it sees to what it has learned. While our YOLO-based object detection model shows promising results, we acknowledge the necessity for further validation with larger and more diverse datasets before advocating for widespread practical implementation.

For this project, the algorithm was trained on a Windows 11 PC equipped with a Core i9 3 GHz CPU, 32 GB of RAM, and a 24GB Nvidia RTX 4090 GPU. Implementation and training were carried out using the PyTorch deep learning library (Paszke *et al.* 2019), with support from the Ultralytics codebase (Jocher *et al.* 2022). The process is easy to set up and could be divided into two steps: (1) Annotating the dataset with respective objects, and (2) Training an ML model on the dataset. A small dataset, say 100–150 images, could be annotated in 2–3 h which in turn could be used to train an ML model (say 1–2 h). In other words, this sort of system could be set up in only 4–5 h by someone comfortable with running desk-based computer models. Significant coding expertise is not required.

To create a training and validation split, the dataset was randomly partitioned into 80% for training and 20% for validation. The training dataset was manually categorized into eight distinct classes, including: 'bottle 0.5,' 'bottle 1,' 'drum 20,' 'jug 5,' 'jug10,' 'bucket 5,' 'bucket 10,' and 'kholash 5'. The numeric suffix in each class label specifies the volume of the respective object. The training set also included some background images designed to help train disambiguation into the model (distinction of foreground, background, etc.). When unknown images are presented to the trained model it places bounding boxes around predicted objects (Figure 2).

Evaluation of results was conducted through the creation of a confusion matrix and statistics on precision (ratio of true positives to all positives) and recall (ratio of true positives to true positives plus false negatives). Model performance was represented through precision (P), recall (R) and mean average precision (mAP) @ 0.5 intersection over union statistics.

HOW GOOD WERE THE MODELS AT IDENTIFYING WATER CONTAINERS?

Of the five ML models assessed, it was YOLOv5s that performed best across all types of water storage containers, achieving an overall precision of 87.7%. Table 1 presents a quantitative analysis of the five different models tested. Notably, the 'Kholash-5' container type achieves the highest precision at nearly 99%. This outcome may be attributed to the fact that 'Kholash'




Figure 2 | Sample model output on test set.

is the most prevalent water container type within the dataset meaning that the ML model had much more training data for this object type. It is moreover a very distinctive and unique water container shape (Figure 1, middle panel). Conversely, the 'Bucket 5' type of container (5 L cylindrical bucket) was only correctly identified 69.7% of the time and with a higher number of false negatives leading also to a lower recall score. Images including a standard scale object (usually a 500- or 1,000-mL PET plastic water bottle) allowed higher rates of precision across all object classes, and we note that both precision and recall were very good for these containers across all container types.

The Confusion Matrix for YOLO V5s (Figure 3) provides more detailed information about exactly where the model had trouble correctly predicting trained objects. Specifically, and as hinted above, 'Bucket-5' and 'Bucket-10' are more frequently confused with each other. Also, note the tendency in some cases for objects to be confused with background. In the former case, it is likely that perspective makes two different-sized versions of the same object more difficult for the model to distinguish, whilst for the latter case lighting levels and shading may be the culprit. In both cases, additional training with images constructed to replicate commonly encountered combinations of lighting/shading and different-sized versions of the same thing will improve model performance for these objects.

Although not uniformly accurate across all container types, YOLO V5s performed well enough to support estimation of domestic water storage capabilities, to which we turn in the next section.

RESULTS OF PILOT: WATER STORAGE IN TEKNAF UPAZILA, BANGLADESH

Figure 4 shows the results for analysis using YOLO V5s of images collected from 288 households during the data collection phase in 2022. It is clear that in these communities, the 5 L kholash is by far the most common water container used by respondent households. From a WASH planning perspective, this is a useful finding for the following reasons:

- Kholashes are usually made of aluminium, making them more robust than plastic vessels, though also more expensive.
- The design usually allows for a tight-fitting plate-type lid, helping to prevent recontamination of collected water, and
- The vessel is better suited than other vessel types for carrying on the head or the hip, possibly reducing the likelihood of musculoskeletal injuries (Geere *et al.* 2010; Venkataramanan *et al.* 2020).

Though far less prevalent, the next three most common water containers are the 500 mL bottle ($n = 118$), the 10-L bucket ($n = 109$) and the 1 L bottle ($n = 95$), which suggests that very many water containers are in fact reused containers originally obtained for other purposes. Especially for the 500 mL and 1 L plastic bottles, this may bring larger health risks related to their construction from PET plastic and the difficulty of properly cleaning them (Ioannidou *et al.* 2016; Staddon & Brewis 2024).

Table 1 | Precision (P), recall (R) and mean average precision (mAP) @ 0.5 intersection over union (IOU) statistics for the five YOLO models tested

Class	Instances	YOLOv5n			YOLOv5s			YOLOv5m			YOLOv5l			YOLOv5x		
		P	R	mAP	P	R	mAP	P	R	mAP	P	R	mAP	P	R	mAP
All	1,191	0.838	0.787	0.836	0.877	0.851	0.867	0.859	0.809	0.86	0.876	0.814	0.864	0.849	0.805	0.8
Bottle 0.5	70	0.849	0.886	0.896	0.898	0.843	0.95	0.819	0.9	0.92	0.94	0.9	0.945	0.855	0.871	0.91
Drum 20	107	0.888	0.813	0.882	0.899	0.911	0.91	0.929	0.86	0.91	0.92	0.86	0.92	0.883	0.85	0.888
Kholash 5	107	0.962	0.956	0.976	0.98	0.967	0.981	0.986	0.962	0.98	0.981	0.962	0.98	0.978	0.954	0.982
Bottle 1	367	0.924	0.787	0.876	0.902	0.898	0.92	0.885	0.806	0.874	0.903	0.861	0.903	0.86	0.815	0.795
Bucket 5	108	0.629	0.423	0.511	0.697	0.635	0.59	0.652	0.51	0.581	0.703	0.433	0.562	0.667	0.462	0.445
Jug 5	104	0.753	0.753	0.806	0.855	0.801	0.88	0.892	0.788	0.873	0.853	0.753	0.834	0.855	0.767	0.58
Jug 10	55	0.941	0.873	0.931	0.935	0.927	0.947	0.923	0.873	0.931	0.943	0.909	0.946	0.924	0.891	0.833
Bucket 10	234	0.761	0.803	0.813	0.847	0.829	0.824	0.785	0.774	0.809	0.768	0.833	0.819	0.767	0.833	0.804

The bold values show the best performance of the given metrics in a particular 'class'.




Figure 3 | The normalized confusion matrix for the YOLO V5s model.




Figure 4 | # Instances of each water container type.

Next, we looked at the mix of water containers used by households to fetch and store water. As noted previously, the kholash is the most prevalent water container, possessed by 66% of households. The most common combination of water containers is one 500 L kholash, and one 10 L bucket, which may well represent the maximum carrying capacity for a single person.¹

Finally, we present the total volume of water storage per household by utilizing model results together with the household population data also collected as part of the survey process. Figure 5 shows that the vast majority of households have significantly less than 10 L of water stored per person, with the mean average storage being 4.5 L/person and the modal (most

¹ The average water fetching time (including waiting) was 30 min per trip, with most households making 1–2 trips per day.




Figure 5 | Frequency of households with specified volumes of per-person water storage.

common) value is 2.5 L/person implying that at least six water fetching trips might be required each day to achieve WHO guidelines for emergency humanitarian response. Since the vast majority of surveyed households reported that they made three or fewer trips per day, it is clear that basic needs are not being met.

CONCLUSION: UTILITY OF ML FOR ASSESSING DOMESTIC WATER STORAGE

In the absence of a piped to home water supply, a key pinch point in achieving SDG 6.1 (Universal Access to Clean and Abundant Drinking Water) is the availability of containers that can safely hold a sufficient volume of water for domestic needs. This paper demonstrates a novel method of extracting useful water storage data from photographs taken during conventional household WASH surveys. Using a readily available ML system called YOLO V5s we were able to extract information about quantities of water stored and storage vessel types with a high degree of accuracy and reliability. With this new data in hand, it was easy to calculate water sufficiency relative to reported household size and compare this against the WHO 15 L per person per day standard. We were also able to further contextualize this storage information with collected data on the number and length of water-fetching journeys per day. We suggest such ML-based techniques could, with only a small amount of training, become a useful tool in water security research and humanitarian practice.

From a technical perspective, we note that even with relatively few training images (less than 700) the ML algorithm demonstrated impressive precision and recall, even with complications related to poor lighting, shading, occlusion of target objects by other objects, angles of view, etc. Notwithstanding these difficulties, we achieved 90% accuracy or better for all but one object class. Further training of the algorithm, with either real-world images or images specially constructed to be challenging, will only improve object detection accuracy.

Our dataset is limited in size as this is a proof-of-concept paper. However, we have employed techniques such as data augmentation/cross-validation to improve generalizability. Additionally, we are exploring the integration of larger datasets and alternative approaches (e.g., ensemble methods) to enhance model robustness in future work.

We also expect that other models, such as Microsoft's Florence-2 may be able to more efficiently automate the image analysis process, including offering summary narratives with desired statistics for each image analysed as well as for entire sets of images.

While we have made progress in using ML techniques to describe water storage quantitatively, this approach may have only limited relevance to water quality assessment. Specifically, while it may be possible to identify those that are excessively dirty or stored on dirt floors these characteristics are only indirectly linked to the quality of stored water.

DATA AVAILABILITY STATEMENT

Data cannot be made publicly available; readers should contact the corresponding author for details.

CONFLICT OF INTEREST

The authors declare there is no conflict.

REFERENCES

- Akhter, M., Uddin, S. M. N., Rafa, N., Hridi, S. M., Staddon, C. & Powell, W. (2020) Drinking water security challenges in Rohingya refugee camps of Cox's bazar, Bangladesh. *Sustainability*, **12** (18), 7325. <https://doi.org/10.3390/su12187325>.
- Chowdhury, P. N., Bhunia, A. K., Sain, A., Koley, S., Xiang, T. & Song, Y.-Z. (2023) What can human sketches do for object detection?, 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Vancouver, BC, Canada: IEEE, 15083–15094. <https://doi.org/10.1109/CVPR52729.2023.01448>.
- Diwan, T., Anirudh, G. & Tembhurne, J. V. (2023) Object detection using YOLO: challenges, architectural successors, datasets and applications, *Multimedia Tools and Applications*, **82** (6), 9243–9275. <https://doi.org/10.1007/s11042-022-13644-y>.
- GBD 2016 Diarrhoeal Disease Collaborators (2018) Estimates of the global, regional, and national morbidity, mortality, and aetiologies of diarrhoea in 195 countries: a systematic analysis for the Global Burden of Disease Study 2016, *The Lancet Infectious Diseases*, **18** (11), 1211–1228. [https://doi.org/10.1016/S1473-3099\(18\)30362-1](https://doi.org/10.1016/S1473-3099(18)30362-1).
- Fan, B. & Song, X. (2024) A review of deep learning based object detection algorithms, *Journal of Engineering Research and Reports*, **26** (11), 88–99. <https://doi.org/10.9734/jerr/2024/v26i111316>.
- Geere, J.-A. L., Hunter, P. R. & Jagals, P. (2010) Domestic water carrying and its implications for health: a review and mixed methods pilot study in Limpopo Province, South Africa, *Environmental Health*, **9** (1), 52. <https://doi.org/10.1186/1476-069X-9-52>.
- Ioannidou, M. D., Samouris, G. & Achilias, D. S. (2016) Acetaldehyde contamination of water, alcoholic, and non-alcoholic beverages stored in glass or plastic bottles, *Toxicological & Environmental Chemistry*, **98** (10), 1183–1190. <https://doi.org/10.1080/02772248.2015.1115505>.
- Jocher, G., Chaurasia, A., Stoken, A., Borovec, J., NanoCode012, Kwon, Y., Michael, K., TaoXie, Fang, J., imyhxy, Lorna, Zeng, Y., Wong, C., Abhiram, V., Montes, D., Wang, Z., Fati, C., Nadar, J., Laughing, UnglvKitDie, Sonck, V., tkianai, yxNONG, Skalski, P., Hogan, A., Nair, D., Strobel, M. & Jain, M. (2022) Ultralytics/yolov5: v7.0 - YOLOv5 SOTA Realtime Instance Segmentation. <https://doi.org/10.5281/ZENODO.7347926>.
- Kashtan, V. Yu. & Hnatushenko, V. V. (2024) Machine learning for automatic extraction of water bodies using Sentinel-2 imagery, *Radio Electronics, Computer Science, Control*, **1**, 118. <https://doi.org/10.15588/1607-3274-2024-1-11>.
- Liang, S., Wang, W., Chen, R., Liu, A., Wu, B., Chang, E.-C., Cao, X. & Tao, D. (2024) Object detectors in the open environment: challenges, solutions, and outlook. arXiv preprint arXiv:2403.16271. <https://arxiv.org/abs/2403.16271>.
- Liu, B. & Li, T. (2024) A machine-learning-based framework for retrieving water quality parameters in urban rivers using UAV hyperspectral images, *Remote Sensing*, **16** (5), 905. <https://doi.org/10.3390/rs16050905>.
- Meierhofer, R., Wietlisbach, B. & Matiko, C. (2019) Influence of container cleanliness, container disinfection with chlorine, and container handling on recontamination of water collected from a water kiosk in a Kenyan slum, *Journal of Water and Health*, **17** (2), 308–317. <https://doi.org/10.2166/wh.2019.282>.
- Nurjana, M. A., Srikanthi, Y., Wijatmiko, T. J., Hidayah, N., Isnawati, R., Octaviani, O., Gunawan, G., Samarang, S., Koraang, M. E., Anastasia, H. & Ningsi, N. (2023) Water containers and the preferable conditions for laying eggs by *Aedes* mosquitoes in Maros Regency, South of Sulawesi, Indonesia, *Journal of Water and Health*, **21** (11), 1741–1746. <https://doi.org/10.2166/wh.2023.270>.
- Opreszko, M. C., Guo, Y., MacDonald, L., MacDonald, L., Kiihl, S. & Schwab, K. J. (2013) Impact of water-vending kiosks and hygiene education on household drinking water quality in rural Ghana, *The American Journal of Tropical Medicine and Hygiene*, **88** (4), 651–660. <https://doi.org/10.4269/ajtmh.12-0065>.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J. & Chintala, S. (2019) PyTorch: an imperative style, high-performance deep learning library. arXiv. <http://arxiv.org/abs/1912.01703> (Accessed: 18 July 2024).
- Prüss-Ustün, A., Wolf, J., Bartram, J., Clasen, T., Cumming, O., Freeman, M. C., Gordon, B., Hunter, P. R., Medlicott, K. & Johnston, R. (2019) Burden of disease from inadequate water, sanitation and hygiene for selected adverse health outcomes: an updated analysis with a focus on low- and middle-income countries, *International Journal of Hygiene and Environmental Health*, **222** (5), 765–777. <https://doi.org/10.1016/j.ijheh.2019.05.004>.
- Rafa, N., Uddin, S. M. N. & Staddon, C. (2020) Exploring challenges in safe water availability and accessibility in preventing COVID-19 in refugee settlements, *Water International*, **45** (7–8), 710–715, DOI: <https://doi.org/10.1080/02508060.2020.1803018>.
- Shahbaz, A. & Jo, K.-H. (2021) Enhanced unsupervised change detector for industrial surveillance systems, *IEEE Transactions on Industrial Electronics*, **68** (9), 8973–8981. <https://doi.org/10.1109/TIE.2020.3013747>.
- Staddon, C. & Brewis, A. (2024) Household Water Containers: mitigating risks for improved Modular, Adaptive, and Decentralized (MAD) water systems, *Water Security*, **21**, 100163. <https://doi.org/10.1016/j.wasec.2023.100163>.
- Venkataramanan, V., Geere, J. A. L., Thomae, B., Stoler, J., Hunter, P. R., Young, S. L., Adams, E., Ahmed, J. F., Alexander, M., Balogun, M., Boivin, M. J., Brewis, A., Carrillo, G., Chapman, K., Cole, S., Collins, S. M., Escobar-Vargas, J., Eini-Zinab, H., Freeman, M. C., Ghattas, H., Ghorbani, M., Hagaman, A., Jamaluddine, Z., Jepson, W. E., Krishnakumar, D., Maes, K., Mathad, J., Maupin, J., Melgar-Quinonez, H., Miller, J. D., Morales, M. M., Moran, J., Omidvar, N., Owuor, P. M., Rasheed, S., Rosinger, A. Y., Samayoa-Figueroa, L., Santoso, M. V., Schuster, R. C., Sheikhi, M., Srivastava, S., Staddon, C., Sullivan, A., Tesfaye, Y., Triviño, N., Trowell, A., Tshala-Katumbay, D., Tutu,

- R. & Wutich, A. (2020) In pursuit of 'safe' water: the burden of personal injury from water fetching in 21 low-income and middle-income countries, *BMJ Global Health*, **5** (10), e003328. <https://doi.org/10.1136/bmjgh-2020-003328>.
- Wang, Z., Li, Y., Chen, X., Lim, S., Torralba, A., Zhao, H. & Wang, S. (2023a) Detecting everything in the open world: towards universal object detection, *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Vancouver, BC, Canada: IEEE, 11433–11443. <https://doi.org/10.1109/CVPR52729.2023.01100>.
- Wang, C.-Y., Bochkovskiy, A. & Liao, H.-Y. M. (2023b) YOLOv7: trainable bag-of-freebies sets new state-of-the-art for real-time object detectors, *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Vancouver, BC, Canada: IEEE, 7464–7475. <https://doi.org/10.1109/CVPR52729.2023.00721>.
- Wani, H., Smeets, P. & Shrivastava, S. (2022) Evaluation of WASH indicators associated with diarrhoeal disease among under-five children in an urban slum pocket, Mumbai city, India: a community-based repeated cross-sectional study, *Journal of Water, Sanitation and Hygiene for Development*, **12** (4), 359–374. <https://doi.org/10.2166/washdev.2022.196>.
- Wu, Q., Shen, C., Wang, P., Dick, A. & van den Hengel, A. (2018) Image captioning and visual question answering based on attributes and external knowledge, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **40** (6), 1367–1381. <https://doi.org/10.1109/TPAMI.2017.2708709>.

First received 12 August 2024; accepted in revised form 1 May 2025. Available online 6 May 2025