

Reviewer 2

The authors present an interesting and novel hardware/software package for computing river streamflow using camera imagery. In their package called the Computer Vision Stream Gauging (CVSG) system, they deploy a stereoscope camera paired with an edge computing device to capture stereo video in order to process into velocities, cross sectional geometry, and water surface plane determination. The CVSG is unique in that it does not require ground control points to perform image velocimetry analysis. Also, the CVSG uses a machine learning technique to improve results over time at a site as the equipment is allowed to collect more measurements. The CVSG seems to perform at least as well as other commercially available image-based velocity software. In one particular case, the CVSG completely outperforms other techniques, in part because it can interpret complex water surface profiles, such as cases with standing waves or other extreme conditions which normally highly degrade other common image velocimetry approaches. The authors do a good job evaluating the performance of the CVSG velocity and discharge results against standard/conventional methods and show excellent agreement. Where the system performs worse than conventional methods, the authors explain the limitations of the system. The authors should consider explaining a bit more about the approach for computing learned water surface plane and cross-sectional shape with the stereoscopic imagery. As with other non-contact streamflow methods, the detection of error in cross-sectional error is difficult. Since the CVSG system can presumably develop information about the cross-sectional geometry owing to the stereoscopic camera approach, the authors may consider expanding the work a bit to describe, if applicable, the performance of the CVSG to determine or evaluate changes in cross-sectional area compared to a priori knowledge (e.g., input surveys of the cross-section). Otherwise, I believe this manuscript describes a useful advance in non-contact streamflow techniques.

Reply: We thank the reviewer for their time and feedback about the manuscript. The authors agree that there is significant untapped potential in the rich amount of spatial data that is generated from the stereoscopic imagery that is computed where the current implementation (as described) only begins to scratch the surface of what is possible. Detail has been added about the simple approach for the site learning cross-section in section 2.2 from L150, however none of the sites or data presented made use of this aspect and this fact has been similarly clarified (all were provided with fixed manually surveyed cross-section profiles that were not free to be adapted to the ongoing live surveying of the terrestrial profile provided by the stereophotogrammetry). We believe this remains an important area for future research and evaluation requiring studies over longer timespans with ongoing manual surveying comparisons. We have added this important note to the discussion section at L681.

Added from L150: “While the CVSG system maintains an adaptive cross-section database for each site which is compared and adapted with each measurement for visible terrain above the estimated water level (applying more weight to gradual changes in time and requiring many consistent measurements to gradually apply any observed dramatic changes in the cross-section profile), the results of this study applied fixed manual cross-section surveys from the time of deployment over the entirety of the time periods evaluated.”

Added from L671: “A comprehensive study evaluating the use of stereo camera systems such as CVSG for quantifying adaptive cross-sections is an important area of future research to be determined over studies spanning longer timescales with significant erosion and/or accretion events at suitable study sites.”

In addition to the comments above, I note several comments below, organized by line number. After consideration, I would happily recommend that the manuscript be accepted with some minor revisions.

Line 65: I agree. In part this is because there have been few available quality software products to aid in adoption by hydrometric agencies. This is starting to change, but it is a slow process.

Reply: We thank the reviewer for sharing their view point on this matter and the authors agree with these comments as it is in agreement with our experience in discussion with various hydrometric agencies.

Line 105: It would be very helpful to report the CVSG power consumption information. I would presume that the camera is a significant portion of the power budget.

Reply: The authors thank the reviewer for this feedback. We have added details of the characteristic power budget of the deployed CVSG hardware at L107.

Added to L107: “The total power consumption of the CVSG hardware collecting data in this study was on the order of 36 W hr per day, averaging 1.5 W with a peak power draw of 30 W.”

Line 118: IMUs are notorious for drifting when position is fixed for a long period of time. Can you describe how the CVSG accounts for this? Do you observe IMU drift? Owing to the requirement of seeing a horizon line, are you correcting IMU drift from the horizon?

Reply: We fortunately have not observed any notable IMU drift from any CVSG hardware installations after the initial factory calibration has been applied. We only use the linear acceleration in the x, y and z dimensions to determine the orientation of the camera for each measurement compared to the last known orientation. We have observed drift in the pose estimated by the same IMU under moving measurement applications, but the absolute measure of the orientation of the camera relative to the force of gravity has proven to be robust (even in moving applications). We do not require or use the horizon line for any corrections, as this would be quite restrictive for the selection of appropriate sites. We have clarified the wording from L139 to reduce any confusion about the horizon in the field of view. This modification is included in the modification to the same text in the next comment.

Line 140: I am presuming that this requirement is to avoid glare? Can you expand on this statement?

Reply: The authors thank the reviewer for this comment. The reviewer’s presumption is correct, and we have added this expansion to L139 to improve the clarity of the reasoning for the guidance provided.

Modified from L139: “When selecting a site, care should be taken to identify sites with suitable surface flow visibility and oriented south-facing (southern hemisphere) or north-facing (northern hemisphere) where possible to avoid sun glare, while keeping the horizon or sky outside of the camera field of view (maximising the water surface in the field of view and reducing automatic exposure determination from the sky).”

Line 141: So this would indicate a high-oblique view? As in the sky is not visible in the field of view? If so, does this eliminate the possibility of IMU calibration using the horizon line?

Reply: This is correct that the sky is preferably not visible in the field of view from the perspective of maximising the field of view available for measurement of the stream surface while minimising any negative effects from autoexposure to the brightness of the sky. We have clarified the wording here

in L139 to reduce the possibility for this misunderstanding of the camera angle recommendation. This modification is included in the modification for the previous comment.

Line 145: Can you provide details about how the stereoscopic camera determines the the land surface? What quality assurance methods are in use? What is the accuracy of the stereoscopic transformation...the parallax of the camera in use is fairly small, so I would expect that there is potentially significant errors in the transformation process. What about obstructed views owing to shadows, obstructions, etc. (e.g. boulders or even cobbles may present "shadows" unseen by the stereo camera ... what are the impacts to cross-section geometry accuracy from these sorts of artifacts?)

Reply: The authors thank the reviewer for this suggestion to provide this detail. We have added this information to this section 2.2 at L151, significantly improving the description that was provided previously. For any obstructed views which are not unobstructed along any of the search lines within the cross-section analysis region, this would simply appear as missing data, and you would seek to avoid setting up the system at a site with obstructions like this without at least partially providing manual survey data covering the obstructed areas.

Added from L151: “Stereophotogrammetry is applied to estimate the distance from the camera to features which are matched between the stereo pairs of rectilinear corrected images where a convolutional neural network model (provided by the camera manufacturer, Stereolabs), that has been trained on pairs of stereo images, is applied to improve both the accuracy and solution density particularly with reflective and featureless surfaces.”

Line 161: Because you are using an IMU and stereo camera to determine the water plane, is there potential that the CVSG would be able to better handle high-slope systems, where the typical mono-lens camera approaches solve for a water surface plane that is oriented parallel to the Z coordinate (which respect to gravity)? This would be a useful differentiation between typical camera-matrix solution approaches for rectification vs epipolar geometric approaches using stereo cameras.

Reply: The authors thank the reviewer for raising this point. We have added this note in a modification from L147.

Modification from L147: “However, a stereo computer vision system also makes it possible to initially survey and then continuously monitor the terrain of the cross-section above the water level for changes due to erosion, deposition, or vegetation, and offers the potential advantage for measuring surface velocities on variable or steep hydraulic gradients.”

Line 183: How is averaging of the flow field stack able to suppress motion artifacts? Wouldn't a median be better?

Reply: We agree in principle that applying a median would be more idealistic for accurate surface velocity estimates across a wide range of conditions relating to optical flow visibility, and early tests trialled both approaches. However, in practise we have opted for the analysis of the presented results to continuously add each frame to an accumulating flow field and take the average by dividing the number of frames in order to reduce the computational hardware requirements for memory usage when scaling up to longer duration measurements consisting of many frames while maintaining compatibility with the existing edge computing hardware. This has now been noted as an addition from L183.

Added from L183: “While taking the median of the flow fields would be reasonably more preferable in this context, the average accumulating flow field computation is applied to reduce the edge computing hardware

requirements of the method, particularly with memory usage as the duration of the measurement scales the number of instantaneous flow field frames stored in memory for a median calculation.”

Line 187: Would would these erroneous motions indicate physically? For example, would camera motion (e.g., wind for example buffeting the instrument) be one of these extraneous motions?

Reply: The authors have observed that the physical erroneous motions that are referred to in this statement relate to wind shaking vegetation or swinging ropes in the foreground of the frame, as well as bugs and animals that are not moving in the assumed plane of the water surface. Most oscillatory motions (such as the result of wind buffeting the instrument) are removed or reduced in the flow field accumulation step explained in a prior step of the methodology. We have added examples of these erroneous motions to L187.

Modification from L187: “From this point, the motions out of the assumed plane of the water surface are filtered out of the analysis to further remove false motions unrelated to the waterway surface velocities (such as animals and swinging ropes which are not moving in the assumed plane of the water surface).”

Line 228: This is an improvement over the standard approach of one alpha value per section. I agree with this approach.

Reply: We thank the reviewer for this comment, as we have found this to be a reliable approach to date.

Line 245: Although I understand the scope of this paper is to compare the CVSG to other commercially available image velocimetry approaches, I think it would strengthen the work to also show results using some of the other well-adopted approaches in the literature, for example Patalano's RIVeR (<https://doi.org/10.1016/j.cageo.2017.07.009>) Perk's KLTIV (<https://doi.org/10.5194/gmd-13-6111-2020>)

Reply: The authors agree that a broad and comprehensive comparison between the intricacies of results and parameters using an array of surface velocity methods would be valuable research, particularly towards the development of an ensemble-based surface velocity analysis. However, the scope of this initial paper focusing on explaining the methodology and evaluation of the CVSG system for adaptive streamflow monitoring is already quite extensive for a single publication.

Figure 2: It would be helpful to annotate these images to include the region of interest used by each method for computation of discharge. For example, your results demonstrate that the CVSG dramatically outperforms other methods for the irrigation channel in NSW. I'd like to know where the STIV search lines were placed? What is the ROI for the SSIV processing?

Reply: The authors thank the reviewer for this feedback, and agree that this is a valuable illustration to provide to improve Figure S1 in the supplementary material with an annotated image. We have added this image as a second panel for Figure S1 (generated through programming a procedural pattern to apply to the coordinates for each region of interest).

Modification of Figure S1:





Figure S1: Raw discharge measurements using different technologies along the length of the irrigation channel in NSW, Australia.

Line 292: Perhaps reword to indicate "diffuse light" rather than "softly lit" -- The light diffusion leads to reductions in shadows, making this dataset a great test for these methods.

Reply: The authors thank the reviewer for this suggestion, and agree that this is better phrasing. We have made this change in line with this suggestion on L292.

Modification of L292: "This benchmark case study presents a favourably diffusely lit environment with visible surface rippling features across the full width of the cross-section, and a sky/vegetation reflective water surface."

Figure 3:

It would be helpful to indicate that the ADCP data are the near-surface cell values. Additionally, label the plot to indicate these are surface velocity profiles. Although this is indicated in the text, it should also be included in the caption and/or figure legend.

Additionally, panel D may benefit from also including a residual plot. It seems that there is visually a trend in the CVSG results of under-predicting lower and over-predicting higher Qs. Alternatively, linear trend-lines could be added to show whether this is the case or not.

Reply: We thank the reviewer for these valuable suggestions for improving the manuscript. We have added the indication of the near-surface cell ADCP observations to both the figure legend and caption, as well as changing the y-axis for the surface velocity profiles to be labelled as such. The authors also thank the reviewer for their suggestion for Figure 3d, however the only significant trend present between lower and higher Qs is in the CVSG discharges applying the stereophotogrammetry estimated water levels, where this trend is already clearly observed and noted in the discharge rating of Figure 7 and L593-596.

Modification of Figure 3:

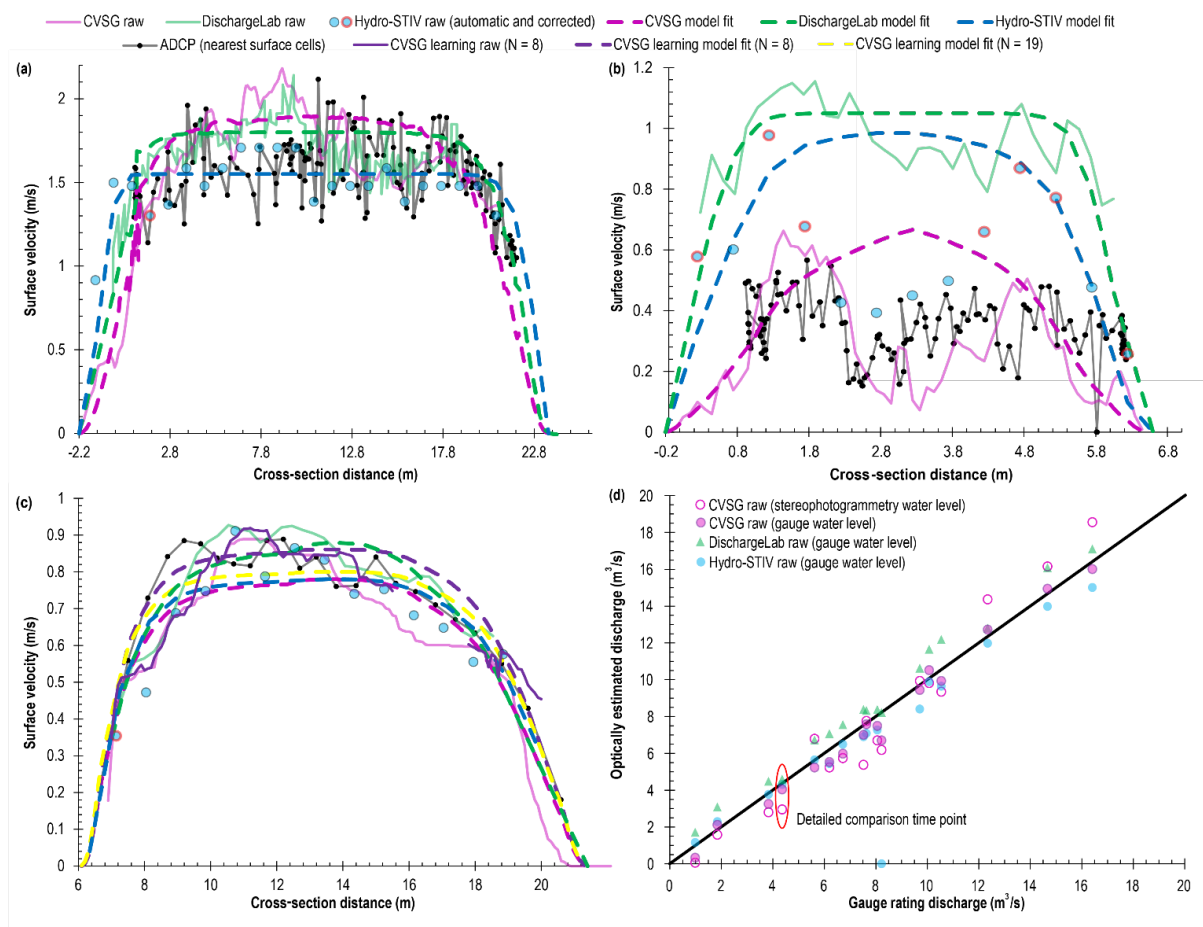


Figure 2: Detailed time point comparison raw and model fitted velocity measurements plotted with nearest surface ADCP measurement cells over the cross-sections at (a) Castor River, Ontario, Canada, (b) an irrigation channel in NSW, Australia, and (c) Tyenna River, Tasmania, Australia. (d) Correlation plot between the gauge rating and optically estimated discharges at comparison time points at Tyenna River, Tasmania, Australia, with the detailed comparison time point indicated. CVSG 5-second duration surface velocities shown for (a) Castor River, Ontario, and (b) the irrigation channel in NSW, Australia. CVSG 10-second duration surface velocities shown for (c, d) Tyenna River, Tasmania, Australia. Hydro-STIV velocity estimates outlined in black were automatically produced, whereas the estimates outlined in red were corrected to the Fourier result or manually corrected to reduce automatically overestimated velocities resulting from the higher frequency surface wave patterns or underestimated tracer-poor search lines.

Line 479: This makes sense, because of potential errors in the determination of WSE from the stereo cameras.

Reply: The authors thank the reviewer for their feedback on the reasonableness of this line.

Line 500: Was this primarily caused by clear water? Relatedly, does the CVSG manage to see any of the bed through the water? Perhaps if so, there might be some value in attempting to extract bed geometry, assuming a suitable correction for refractive properties can be found.

Reply: This is precisely our understanding of the data. We have previously experimented and found surprising success with the capability to use refractive correction in the reconstruction of the submerged cross-section under the clear water conditions at the Tyenna River site, but have yet to apply and evaluate this approach across further sites or develop automated logic for when this is suitable to apply, and how to adaptively integrate these submerged cross-section measurements into the site cross-section.

Line 561: The ability to capture low flow (high clarity water) image velocimetry measurements accurately continues to be a substantial challenge.

Reply: The authors agree with this comment by the reviewer, and welcome future research towards non-contact methodologies that specifically improve measurements under these conditions.

Line 674: Based on the findings of this paper, I agree with the later concept that CVSG can help identify low-flow site suitability. I am less convinced that CVSG (or any other image velocimetry approach) will inform low flow conditions.

Reply: We agree with this assessment of image velocimetry approaches in general, and share this experience. L674 does not seek to make any statement that CVSG can inform low flow conditions, rather this line describes a practical solution to this limitation by providing the lower flow ratings to CVSG through manual gaugings where flow conditions are generally also safer for personnel and equipment.

Section 5: Conclusions: One thing not discussed in this paper is the errors associated with cross-sectional area. It seems that since the CVSG is able to extract stereographic elevations of the low flow channel (or better yet dry channel), that there should be a way to consider changes in cross-sectional area ratings. Maybe something for a future paper?

Reply: The authors thank the reviewer for this recommendation. We have added this to L736 in the conclusions.

Added from L697: “This work did not address errors associated with cross-sectional area changes and the capability of the CVSG system to extract stereophotogrammetry estimated elevations of the dry channel areas to inform changes to discharge ratings, which is recommended for future research using stereo imagery-based optical stream gauging approaches.”

Line 706: add a period to the end of the code availability sentence.

Reply: We thank the reviewer for highlighting this oversight. We have added a period as suggested.

Added to L706: “Code not available at the time of publication.”

In addition to my official refereed comments, I would express agreement with RC1's general comments. Although I share RC1's frustration with the lack of explicit and thoroughly detailing the CVSG algorithms. I think that there is a valid criticism here that the paper results would be hard to repeat or evaluate given the details provided.

If the authors can expand on the details such as to address RC1's comments, I believe that many of my original comments will also be addressed.

Reply: We thank the reviewer for adding their feedback on the clarity in the repeatability of the methods detailed in the manuscript. The authors have made changes to the manuscript based on the valuable feedback provided by the reviewers towards clarifying all of the points of confusion surrounding how the methods have been developed and applied.