SHORT COMMUNICATION

A fully automated method for late ventricular diastole frame selection in post-dive echocardiography without ECG gating

Eric Markley [1], David Q. Le [1], Peter Germonpré [2,3], Costantino Balestra [3,4], Frauke Tillmans [5], Petar J. Denoble [5], John J. Freiberger [6], Richard E. Moon [6], Paul A. Dayton [1], Virginie Papadopoulou [1]

[1] Joint Department of Biomedical Engineering, The University of North Carolina at Chapel Hill and North Carolina State University, USA

[2] Center for Hyperbaric Oxygen Therapy, Military Hospital Brussels, Belgium

[3] Environmental & Occupational (Integrative) Physiology Laboratory, Haute Ecole Bruxelles-Brabant (HE2B), Brussels, Belgium

[4] European Research Division, Divers Alert Network, Belgium

[5] Research Division, Divers Alert Network, U.S.

[6] Center for Hyperbaric Medicine and Environmental Physiology, Duke University, U.S.

CORRESPONDING AUTHOR: Virginie Papadopoulou, PhD – *papadopoulou@unc.edu* or *virginiepapadopoulou@gmail.com*

ABSTRACT

Venous gas emboli (VGE) are often quantified as a marker of decompression stress on echocardiograms. Bubble-counting has been proposed as an easy to learn method, but remains time-consuming, rendering large dataset analysis impractical. Computer automation of VGE counting following this method has therefore been suggested as a means to eliminate rater bias and save time. A necessary step for this automation relies on the selection of a frame during late ventricular diastole (LVD) for each cardiac cycle of the recording. Since electrocardiograms (ECG) are not always recorded in field experiments, here we propose a fully automated method for LVD frame selection based on regional intensity minimization. The algorithm is tested on 20 previously acquired echocardiography recordings (from the original bubble-counting publication), half of which were acquired at rest (*Rest*) and the other half after leg flexions (*Flex*). From the 7,140 frames analyzed, sensitivity was found to be 0.913 [95% CI: 0.875-0.940] and specificity 0.997 [95% CI: 0.996-0.998]. The method's performance is also compared to that of random chance selection and found to perform significantly better (p<0.0001). No trend in algorithm performance was found with respect to VGE counts, and no significant difference was found between *Flex* and *Rest* (p>0.05). In conclusion, full automation of LVD frame selection for the purpose of bubble counting in post-dive echocardiography has been established with excellent accuracy, although we caution that high quality acquisitions remain paramount in retaining high reliability.

INTRODUCTION

Venous gas emboli (VGE) can be detected using 2D echocardiography post-dive and are often used as a metric of decompression stress in diving research [1]. It is well established that higher VGE grades are associated with a higher probability of decompression sickness (DCS), but this is not a one-to-one relationship, and VGE are an imperfect surrogate endpoint for DCS [2,3]. Previous work showed that VGE counts, where VGE are counted in a frame with fully open tricuspid valves for each of 10 consecutive heart cycles and averaged, lowers inter-rater variability in comparison to conventional grading for non-medically trained and relatively inexperienced raters [4]. However, bubble counting is a time-consuming process, making it impractical for large datasets and becomes particularly problematic when evaluating VGE dynamics post-dive [5,6]. These often require frequent measurements from the moment the dive ends to more than two hours post-dive and can result in hundreds of datasets for a single experiment. Computer-automation of VGE counting is therefore attractive to both standardize and expedite VGE analysis in diving research.

VGE evaluation as proposed by Germonpré et al. [4] lends itself particularly well to computer automation because VGE are counted only on open-valve frames. The valves can appear discontinuous and difficult to distinguish from VGE when not fully open (as opposed

to end-diastole/proto-systole on apical four-chamber view echocardiography, where they are almost invisible). The chance of misclassifying parts of the cardiac valves as VGE is therefore minimized, and it is relatively easy to count VGE as bright spots against the dark background of the blood-filled chambers in open-valve frames. The time-consuming task of labeling VGE to train machine-learning algorithms can therefore be distributed among many minimally-trained raters.

To fully automate this counting method, selection of a frame during late ventricular diastole (LVD), when the valves are open, is required in each cardiac cycle of the recording. Electrocardiography (ECG) signals can be used for this purpose: either directly during gated acquisition by setting frames to only be saved during this predetermined part of the cardiac cycle, or in post-processing by using the ECG trace to select the right frame. Nevertheless, directly gated acquisitions would discard valuable VGE motion data which help identify VGE, and in any case ECG data is not always recorded in diving field experiments. Depending on ECG brand, model, and settings, ECG systems and algorithms can also introduce time lags in the ECG recordings which can make them more difficult to precisely synchronize with echocardiograph recordings. Therefore having an automated, and ECG-independent LVD frame-selection method remains of interest.

Here we propose a method to select LVD frames on post-dive echocardiograms based on the post-processing of the echocardiogram, using changes in average pixel intensity inside an automatically selected and fixed region of interest (ROI) corresponding to the right atrium. The method's performance is evaluated on previously acquired post-dive echocardiography recordings (complete anonymized dataset with recordings both at rest and after leg flexions of the original bubble counting publication [4]). The basic premise of our method relies on the different echogenicity between blood (dark) and tissue (brighter) on ultrasound imaging. When the right atrium contracts forcing blood out, the fixed ROI on the screen now contains more tissue signal and its average intensity increases. The periodicity of the contraction and expansion of the heart is therefore reflected in peaks and valleys in average intensity, and this is used to locate cardiac cycles and proposed LVD frame selection in the recording.

METHODS
Algorithm for LVD frame selection
A program was developed in MATLAB (R2019b) (The MathWorks Inc., Natick, Massachusetts, U.S.) to automate the selection of a frame during LVD in each cardiac cycle of a recorded post-dive echocardiogram. The different processing steps for this LVD frame selection algorithm are summarized in Figure 1, and then detailed hereafter. A copy of the program is available from the corresponding author upon reasonable request.

First, the user is prompted to select the video to be analyzed. Once selected, it is converted from RGB to grayscale so that its pixel intensity values (which range from 0, completely black, to 255, completely white) are stored in a three-dimensional array with dimensions equal to the number of pixels in the lateral and axial directions for each frame and the number of frames in the video. The average pixel intensity is calculated for each frame, and the average intensity time-series signal is smoothed (moving average filter of width equal to a quarter of the frame rate of the video, rounded up to the closest integer value). The maximum of this smoothed average-frame-intensity-over-time signal is found and the corresponding frame index stored. This chosen frame F_{max}, an example of which is shown in Figure 2A, is then used to estimate the center of the right atrium along with its radius. F_{max} provides a good initial guess to select a frame with closed heart valves, since that frame will mostly show cardiac tissue (brighter than blood, which appears dark on ultrasound). When the valves are closed, the right atrium is the most circular object in the heart, and we can take advantage of this to estimate its radius and center location.

In order to determine the radius and center of the right atrium, F_{max} is up-sampled by a factor of two in each direction using bicubic interpolation. Then F_{max} is blurred using a Gaussian filter. This allows for subpixel resolution when estimating the center of the right atrium. Next, Sobel edge detection is applied to the frame as a necessary precursor to performing a circular Hough transform (CHT). A mask of the same shape as the conic ultrasound view is applied to the edge-detected image to prevent the edge between the black background and the ultrasound image from being detected. This result is seen in Figure 2B. After edge detection has been applied, a custom-written CHT function is performed in order to estimate the center and radius of the right atrium. Since MATLAB's standard CHT function will not detect the right atrium due to it not being circular enough, a custom function for the CHT was written. This
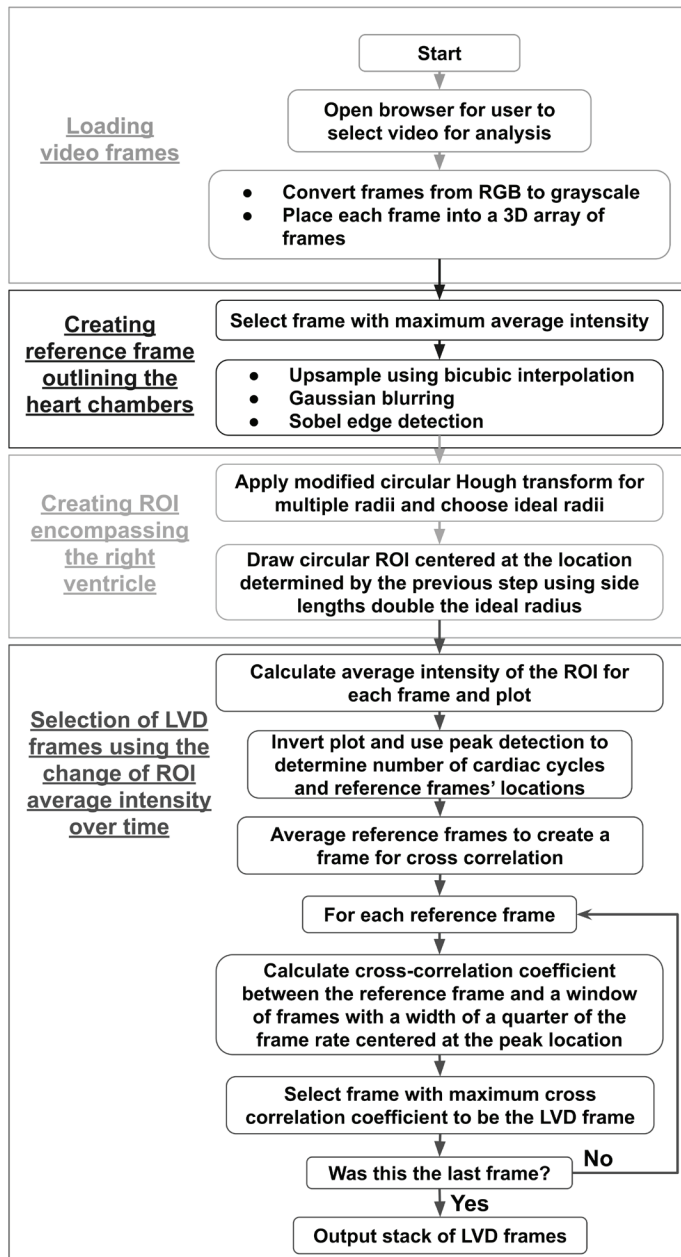
Figure 1
Flowchart describing the key processing steps of the
late ventricular diastole (LVD) frame selection
algorithm implemented in MATLAB.
*Abbreviations: LVD – late ventricular diastole;
RGB: red green blue; ROI: region of interest.*

requested. Each pixel found by the function has its value increased by one over the number of points requested. The final result is blurred, and the maximum value pixel is taken. This is deemed the center of the most circular part of the image for the selected radius. The transform is performed for multiple radii, an example of which is shown in Figure 2C. The overall maximum is taken to be the center of the right atrium, and its associated radius is taken to be the radius of the atrium.

Next, a ROI is selected in each frame. The ROI is a circle centered at the point estimated to be the center of the right atrium. The radius of the ROI is 10 percent larger than the estimated radius of the right atrium. This is done to allow the ROI to still encompass the atrium as the heart beats since the estimated radius using the CHT above comes from F_{max} assumed to be a frame where the atrium is contracted. The average intensity of all pixels inside the ROI is calculated for each frame using a circular mask of the same dimensions [7]. Again, a moving average filter with width equal to a quarter of the frame rate is applied to this time-series signal to act as a low pass filter that discards any frequencies much greater than the heart rate. Peak detection is then performed on this averaged signal to estimate the number of cardiac cycles. Next, a moving average filter with a width of three frames is applied to the original signal to act as a low pass filter. This much narrower filter allows for higher frequencies to pass through compared to the filter used in the previous step. This minimizes the perturbation of the signal during filtering while minimizing noise, ultimately improving the robustness of the peak detection. The signal is then inverted, and peak detection is performed as seen in Figure 3 (the signal must be inverted because MATLAB does not currently support valley detection). Peaks now represent the frames with the ROI of the lowest average intensity for

function takes a radius and a number of points per edge pixel and finds the center of the most circular feature in the image of the given radius. The function finds evenly spaced pixels at a distance equal to the requested radius from each edge-detected pixel based on the number points

Figure 2

Illustration of the key steps to select a ROI from one recording showing:

A) the frame of maximum intensity, F_{max}, used by the algorithm for the determination of right atrium location and radius;

B) the maximum average intensity frame with Sobel edge detection applied (edges thickened for display);

C) an example of the modified CHT applied to the edge detected frame B (overlaid in yellow for reference), here showing the CHT for the final radius chosen (as radii values near the size of the right atrium, the intensity in the center of atrium increases because the circles overlap in the same area);

D) the final ROI selection overlaid on F_{max}.

each cardiac cycle. Therefore, the location of the peaks corresponds to a rough estimate of the frames that are in LVD. Nevertheless, sometimes the actual LVD frame is a few frames before or a few frames after. By averaging the rough estimate of LVD frames (peaks in Figure 3) we get a "reference frame" for what an LVD frame looks like. Then for each detected peak location we determine the cross-correlation coefficient between the averaged reference frame and a window of frames with a width one quarter the frame rate centered at the peak location. The frame with the highest correlation coefficient is taken to be the frame in LVD during that cardiac cycle. Once this has been performed for each peak location, the frames in LVD are output. These frames should all be LVD frames as opposed to a frame with the heart valves closed.

ROI inverted average intensity vs. frame

Figure 3
Inverted average of ROI pixel intensity over each frame in the selected video (blue line),
and after low-pass filtering (black dashed line). Peak detection is shown (red triangles)
and an example frame of a local minimum (close valve frame in **B**) and maximum
(open valve frame in **A**) displayed. Peaks are averaged to create a "reference frame"
(**C**) used to finalize the selection of each LVD frame by maximizing the cross-correlation
value between the reference frame and the frames around each detected peak.

Algorithm performance assessment

The algorithm was tested on all 20 echocardiography videos from the Germonpré et al. bubble counting paper, acquired on with GE 3S-RS sector array probe on a Vivid-i ultrasound system (GE Healthcare, UK) at 30 frames per second [4]. These data are all post-dive acquisitions containing VGE and 10 are recorded at rest (*Rest*) while the other 10 were recorded after leg flexions (*Flex*), both types of recording being routinely used in diving research. MATLAB was used to extract each individual frame of each echocardiogram into a separate folder to allow frame-by-frame scrolling. A human rater (author VP) identified the frame index corresponding to the start and end of each cardiac cycle (also visually confirmed on the ECG trace on each recording) and assessed algorithm-labeled LVD frames

for accuracy in the following manner. Each LVD frame identified by the algorithm for each recording was then evaluated by looking at whether the tricuspid valves were fully open, but also by situating its corresponding frame index with respect to all labeled cardiac cycles to ensure that one and only one LVD per cardiac cycle was identified. Consequently, the following cases were possible for each cardiac cycle (of total frames *n*) of each echocardiogram:

(a) The algorithm picked only one LVD frame within this cycle and it was indeed an open-valve frame: this 1 frame is a true positive and the rest of the frames (*n*-1) are true negatives;

(b) The algorithm did not identify an LVD frame within this cycle: 1 false negative (since one LVD should have been identified) and the rest of the frames (*n*-1) are true negatives;

(c) The algorithm picked only 1 LVD frame within this cycle, but the valves were not fully open: one false positive (the frame picked should not have been selected), 1 false negative (one other frame of the cycle should have been picked as an LVD) and another n-2 true negatives (the remainder of the frames are true negatives); or

(d) The algorithm picked m LVD frames within this cycle, where m>1:

    a. At least 1 of the m picks was indeed an open-valve frame: one true positive, (m-1) false-positives (since only one LVD should be identified per cardiac cycle) and the remainder (n-m) frames are true negatives

    b. None of the m picks were an open-valve frame: m false-positives, 1 false negative (since one LVD should be identified per cardiac cycle) and the remainder (n-m-1) frames are true negatives.

The algorithm's performance was then quantified using the following methods. Sensitivity (equation 1) and specificity (equation 2) in picking LVD frames were calculated using the total number of true negatives, true positives, false negatives and false positives found. Their associated 95% confidence intervals are also computed using the Wilson-Brown method for calculating confidence intervals for proportions in Prism 8 (GraphPad Software, Inc., La Jolla, California, U.S.).

$$Sensitivity = \frac{True\ positives}{True\ positives + False\ negatives} \quad (equation\ 1)$$

$$Specificity = \frac{True\ negatives}{True\ negatives + False\ positives} \quad (equation\ 2)$$

Note that the bubble counting method requires only one "open-valve" frame (LVD) for each cardiac cycle, but there are typically multiple frames during each cardiac cycle that are valid possible picks (in our experience with echocardiography recorded at 30 frames per second this is about four frames in each cycle of about 25 frames). However, once one LVD frame is chosen for each cardiac cycle, all remaining frames pertaining to the same cardiac cycle need to be rejected by the algorithm. The sensitivity and specificity alone cannot capture this, so we also calculated the probability of randomly choosing LVD frames to compare our algorithm's performance to chance, as detailed below.

The probability p_{rec} of randomly choosing c frames for a given echocardiogram containing N total frames, so that each pick is an LVD frame from a different cardiac cycle is:

$$p_{rec} = \frac{c! \times (N\text{-}c)! \times l^c}{N!} \quad (equation\ 3)$$

where l is the number of valid possible LVD in each cardiac cycle. A simple way to compare our algorithm's performance to chance is therefore to compare this with p_{rec} over 20 trials. To do so, we calculate the cumulative binomial probability $P\ (X \geq x)$ with x being the number of echocardiograms for which our algorithm correctly picks one and only one LVD frame in each cardiac cycle without mistakes (out of 20 trials since we have 20 videos in total) and set p<0.05 set as the threshold for significance.

Finally, we assessed whether algorithm performance varied between *Rest* and *Flex* recordings, or depended on the amount of VGE detectable in each recording. Statistical analyses were performed in Prism 8 (Graph-Pad Software, Inc., La Jolla, California, U.S.) and statistical significance levels were set a priori at p<0.05. The specificity and sensitivity were calculated separately for each echocardiogram in the same manner as previously described. Specificities between the *Rest* and *Flex* groups were compared using a Mann-Whitney U test after negative normality test, and the same procedure was repeated for sensitivity comparison. The relationship between VGE counts and algorithm performance was assessed by calculating the correlation between VGE counts and sensitivity, or specificity, respectively. Spearman correlation coefficients (r) and p-values are reported. Data are presented as mean ± standard deviation unless otherwise stated.

RESULTS

The automated LVD frame selection algorithm was used on all 20 echocardiography videos (7,140 frames and 236 cardiac cycles in total). Recordings contained an average of 357 frames and 14.8 cardiac cycles (each consisting of 24 frames on average). From the 7,140 frames analyzed, sensitivity was found to be 0.913 [95% CI: 0.875-0.940] and specificity 0.997 [95% CI: 0.996-0.998], with true positive, true negative, false positive and false negative totals detailed in Table 1.

Using the average number of frames (N=357), cycles (c=15) and overestimating the number of possible correct LVD picks to be half of each cardiac cycle (l=12 out of an average 24 frames per cycle), equation (3) gives p_{rec}=2.0 ×10^{-9}. From this, we find that the probability of observing six or more echocardiograms for which the al-

Table 1

	identified as 'LVD'	identified as 'other'	total
LVD frames	271 (true positives)	20 (false-negatives)	291
other frames	26 (false-positives)	6823 (true negatives)	6849
total	297	6843	7140

Summary of the algorithm's classification results across all 20 videos against human assessment

gorithm correctly picks all LVD frames from the different cardiac cycles without mistakes (as our algorithm did) being due to chance, cumulative binomial probability $P(X \geq 6)$, is p<0.00001. Thus, our algorithm performs significantly better than chance. Note that this estimated chance performance is grossly overestimated already since in reality the number of correct picks per cardiac cycle would be less than half, and it is calculated with the a priori knowledge of the total number of cardiac cycles in the recording which the algorithm does not have (we set 15 picks, whereas the algorithm is not restricted to picking only as many frames as there are cardiac cycles, it calculates how many cycles it thinks a recording has).

The specificity was 0.997 ± 0.003 for *Rest* recordings and 0.997 ± 0.006 for *Flex* recordings. The sensitivity was 0.947 ± 0.103 for *Rest* recordings and 0.881 ± 0.171 for *Flex* recordings. No significant difference was found between *Flex* and *Rest* for either specificity or sensitivity (both p>0.05) and algorithm performance did not depend on VGE counts (r=-0.028 with p=0.91 for sensitivity, and r=0.269 with p=0.25 for specificity).

DISCUSSION

The LVD frame selection algorithm proposed here does not rely on any user input and, as such, is not prone to user bias. It is fully automated from the echocardiography data alone (without necessitating an ECG trace) and could therefore be implemented on data from any ultrasound machine. This is important since not all field research experiments include ECG recordings currently. Since the Germonpré et al. counting method for VGE assessment selects LVD frames for VGE counting, this automated frame selection algorithm is a necessary component toward the full automation of the method. This in turn would allow fast and unbiased automated analysis of large datasets of post-dive echocardiograms. The LVD algorithm takes about eight seconds to run on average (Windows 10.0.18362 Build 18362, Intel(R) Core(TM) i5-3210M CPU @ 2.50GHz, 2501 MHz, 2 Core(s), 4 Logical Processor(s), 16 GB RAM).

Overall, the algorithm seemed to perform well for the purpose of LVD frame identification in the context of VGE counting. It was significantly better than chance at correctly identifying cardiac cycles to pick only one LVD frame for each. Over the 236 cardiac cycles tested in total, specificity and sensitivity exceeded 0.9, and averages from *Rest* and *Flex* recordings separately did not differ and remained high (at or above 0.88). The algorithm's performance did not depend on the amount of VGE present.

In the 20 post-dive echocardiograms used for validation, VGE ranged from 0 to 11.2 (mean 3.38), suggesting that the algorithm is robust over this wide range (grade 0 to grade 4 on the Eftedal and Brubakk scale). It is, however, possible that our intensity-based method would have more difficulty with grade 5 recording since the high number of VGE may add to the chamber brightness (otherwise dark) and influence the ability to correctly estimate the cardiac cycles.

The algorithm relies on peak detection to identify possible LVD frames from each cardiac cycle. This is inherently more difficult when a full cycle is not present, as is sometimes the case in the first and last cycle of a recording. It is therefore interesting to note that from the 16/236 cycles that it failed to choose a frame for (missed those cycles, case (b) described in the methods), six were from either the first or last cycle of a video. Additionally, it misidentified an LVD frame in a 10/236 cycles, of which three were again from the first or last cardiac cycle of a video. This could be mitigated in the future by purposefully ignoring the first and last peak detections for the purposes of bubble counting (the bubble counting method as described in [4] requires 10 consecutive heart cycles).

Although the recordings used in this study used a 30-Hz frame rate, the algorithm is adaptive to the frame rate by design and performance should remain unchanged as long as the frame rate is high enough to adequately sample the heart rate, while retaining at least one LVD frame per cardiac cycle and robustness to acquisition issues. To investigate this, we down-sampled the frames of all 20 videos and found no difference in sensitivity or specificity at 15 Hz compared to 30 Hz (p>0.05), but a significant difference in both sensitivity and specificity at

7.5 Hz compared to 30 Hz (both p<0.0001) (Friedman test for paired video comparison after negative normality test, with Dunn's multiple comparison post-test). Echocardiograms recorded at or above 15 frames per second would therefore be adequately analyzed and, in practice, the frame rate for echocardiography is always above those values clinically.

Finally, it is important to note that the algorithm is particularly sensitive to motion of the heart both in, and out, of the imaging plane. This is because the ROI is selected based on a single frame. This could be corrected by performing the ROI selection process on each frame in the video, but at the cost of a substantial increase in the processing time without parallelization of the code. *Flex* recordings in particular are more difficult to acquire without motion artifacts and sensitivity was lower for those even though the trend was not significant. A crucial aspect for accurate LVD selection (and VGE counting), whether by a human rater or algorithm, remains echocardiography acquisition quality.

CONCLUSIONS

In conclusion, we have developed a fully-automated algorithm for the selection of LVD frames from post-dive 2D echocardiograms, as part of our larger aim to fully automate VGE counting. The algorithm takes advantage of the relatively circular shape of the right atrium to define a ROI around it. Changes in the ROI's average pixel intensity due to the periodic contraction and relaxation of the atrium are then used to identify cardiac cycles, create a reference open-valve frame and propose LVD frames for each cardiac cycle by comparing candidates to that reference. The algorithm performs well on the 20 post-dive echocardiograms tested (total 7,140 frames and 236 cardiac cycles), although we caution that high-quality acquisitions remain paramount in retaining high accuracy.

REFERENCES

1. Mollerlokken A, Blogg SL, Doolette DJ, Nishi RY, Pollock NW. Consensus guidelines for the use of ultrasound for diving research. Diving Hyperb Med. 2016 Mar;46(1):26-32.

2. Doolette DJ. Venous gas emboli detected by two-dimensional echocardiography are an imperfect surrogate endpoint for decompression sickness. Diving Hyperb Med. 2016 46(1):4-10.

3. Nishi RY, Brubakk AO, Eftedal OS. Bubble detection, in: A.O. Brubakk, T.S. Neuman (Eds.) Bennett and Elliott's physiology and medicine of diving, WB Saunders, Philadelphia, PA, 2003, 501-529.

4. Germonpre P, Papadopoulou V, Hemelryck W, et al. The use of portable 2D echocardiography and 'frame-based' bubble counting as a tool to evaluate diving decompression stress. Diving Hyperb Med. 2014; 44(1):5-13.

5. Papadopoulou V, Tillmans F, P. Denoble P. Call for a multi-center study on the intra-subject variability of venous gas emboli, Undersea Hyperb Med. Sept-Oct 2017; 44(5) 377.

6. Papadopoulou V, Germonpre P, Cosgrove D, et al. Variability in circulating gas emboli after a same scuba diving exposure, Eur. J. Appl. Physiol. 2018 118:1255-1264.

7. Shoelson B. createCirclesMask.m (https://www.mathworks.com/matlabcentral/fileexchange/47905-createcirclesmask-m), MATLAB Central File Exchange. Retrieved February 21, 2020.

✦

Markley E, Le DQ, Germonpré P, et al.