IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received December 10, 2019, accepted January 2, 2020, date of publication January 9, 2020, date of current version January 17, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.2965316

Shape Classification Using Hydrodynamic Detection via a Sparse Large-Scale 2D-Sensitive Artificial Lateral Line

BEN J. WOLF[1], PRIMOZ PIRIH[2], MAARJA KRUUSMAA[3], AND SIETSE M. VAN NETTEN[1]

[1]Bernoulli Institute of Mathematics, Computer Science and Artificial Intelligence, University of Groningen, 9747 Groningen, The Netherlands
[2]Biotechnical Faculty, University of Ljubljana, 1000 Ljubljana, Slovenia
[3]Center of Biorobotics, Tallinn University of Technology, 12618 Tallinn, Estonia

Corresponding authors: Ben J. Wolf (ben@benwolf.nl) and Sietse M. van Netten (s.m.van.netten@rug.nl)

This work was supported by the Lakhsmi Project that has received funding from the European Union's Horizon 2020 Research and Innovation Program under Grant 635568.

ABSTRACT Artificial lateral lines are fluid flow sensor arrays, bio-inspired by the fish lateral line organ, that measure a local hydrodynamic environment. These arrays are used to detect objects in water, without relying on light, sound, or on an active beacon. This passive sensing method, called hydrodynamic imaging, is complementary to sonar and vision systems and is suitable for collision avoidance and near-field covert sensing. This sensing method has so far been demonstrated on a biological scale from several to tens of centimeters. Here, we present measurements using a large-scale artificial lateral line of 3.5 meters, consisting of eight all-optical 2D-sensitive flow sensors. We measure the fluid flow as produced by the motion of five different objects, towed across a swimming pool. This results in repeatable stimuli, whose measurements demonstrate a complementary aspect of 2D-sensing. These measurements are both used for constructing temporal hydrodynamic signatures, which reflect the object's shape, and for flow-feature based near-field object classification. For the latter, we present a location-invariant feature extraction method which, using an Extreme Learning Machine neural network, results in a classification F1-score up to 98.6% with selected flow features. We find that, compared to the traditional sensing dimension parallel to the sensor array, the novel transverse fluid velocity component bears more information about the object shape. The classification of objects via hydrodynamic imaging thus benefits from 2D-sensing and can be scaled up to a supra biological scale of several meters.

INDEX TERMS Hydrodynamic imaging, artificial lateral line, neural network, inverse problem, sensor array, feature selection.

I. INTRODUCTION

The lateral line is a fluid flow sensing organ found in fish, that augments their perception [1]. This organ consists of arrays of fluid flow sensors called neuromasts, located on the skin (superficial neuromasts) or in sub-dermal canals (canal neuromasts), each tailored to sensing different properties of the flow [2]. These neuromasts allow fish to sense and act on their local near-field flow situation. They use this sensation for behaviors such as schooling, and detection of both prey and predators [3]. This biological near-field flow sense is often referred to as "touch at a distance" [1]. It does not rely on ambient light and is therefore more reliable than vision in dark or murky environments.

The lateral line has been used as inspiration to create several biomimetic sensor arrays, called artificial lateral lines (ALLs) [4]. The ALL can be either used as a stationary sensor array or, similar to the biological lateral line, attached to a moving vessel, depending on the intended detection task. In both cases, the ALL measures the hydrodynamic effect of relative fluid motion. A stationary ALL can be used for object detection and tracking vibrating dipole sources [5]–[9] or moving objects [10]–[12], with potential use cases including tracking nearby moving vessels. An ALL attached to a vessel can be used for aligning with the freestream flow [13], [14], or for obstacle avoidance [15], [16], potentially enabling

The associate editor coordinating the review of this manuscript and approving it for publication was Haiyong Zheng.

safe navigation for ships or autonomous underwater vehicles (AUVs) [14], [17].

While far-field detection methods such as active sonar can detect ships and other objects, and can resolve their shape over large distances, they rely on an active emitter. Near-field hydrodynamic imaging only works in a short range, but can be more accurate in this range. Furthermore, it is completely passive, enabling covert sensing, and is unaffected by murky waters, an advantage over camera based and sound based solutions. ALL hydrodynamic imaging thus provides a complementary detection method to vision-based and sonar-based systems [5].

Identifying up to three object shapes with an ALL has been achieved via the wake of an object in controlled flow conditions [18]–[22]. There the object is usually placed upstream relative to a sensor array. The resulting wake is repetitive in nature and its periodic features are used for determining the size and shape of objects. An alternative approach [23] assumes still water, where an object moves past the array and creates its own near-field flow. We use this second approach to classify five different objects via their self-produced near-field flow.

The typical length of an ALL ranges from several to tens of centimeters, matching the biological scale of the fish lateral line [4]. The present work describes the first use of a large-scale ALL of several meters and its consequences and adaptations to fluid flow sensor array signal processing. We use eight optical 2D-sensitive flow sensors [24] in a 3.5 m array, to measure the object-produced near-field flow.

We identify key features of the fluid flow for processing the ALL measurements that make object shape classification more robust and less prone to overfitting. There are several automatic feature selection algorithms that can be used for this purpose [25]. Specifically, we employ manual feature selection, filter methods, and wrapper feature selection to determine subsets of features. Then, we use an Extreme Learning Machine (ELM) neural network architecture [26] with cross validation that uses these selected features as input for the final classification step. We also discuss the consequences of sparse spatial sampling for constructing hydrodynamic signatures of shapes, which can be used for further analysis and interpretation.

In the next section, we present background information on hydrodynamic imaging and the ELM neural network. In section III, we describe the experimental setup and the processing pipeline. The fourth section lists the results, which are further discussed and interpreted in section V. We end with a conclusion section, summarizing our main findings.

II. BACKGROUND

A. HYDRODYNAMIC IMAGING

Hydrodynamic imaging [1], [5], [27] is a bio-inspired method which usually applies to sensor arrays that measure a projection of the local hydrodynamic environment in response to a flow source or moving object, see e.g. Fig. 1. In fish lateral line research, the excitation of neuromasts along the trunk, concatenating to a spatial excitation pattern, have been shown to encode the pressure gradient along said trunk [2], [7], [10]. Some fish species have been shown to be capable of detecting objects via hydrodynamic imaging up to a distance of roughly one body length away [7], [28], [29].

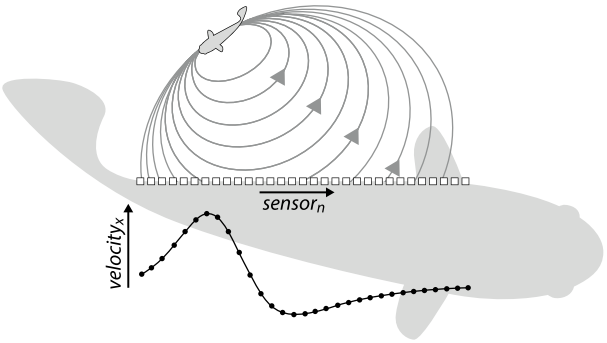

FIGURE 1. Passive hydrodynamic near-field imaging by the fish lateral line. The smaller fish creates a flow field which is sensed by the sensors on the larger fish's body. The instantaneous 1D spatial velocity profile encodes the source's relative location and other properties. Adapted with permission from [30].

An ALL is usually composed of uniformly spaced pressure or fluid velocity sensors [4], sampling the hydrodynamic environment at discrete locations. The concatenation of these discrete samples forms a spatial velocity profile, similar to the spatial excitation patterns measured along the trunk of the fish.

1) VELOCITY PROFILES

With a simplified hydrodynamic model, assuming inviscid flow [31], [32] as described in [7], [10], [11], [33], we can estimate the velocity profile as sampled by an array for a moving or vibrating object; this is called the forward problem. This velocity profile encodes state information such as the relative location of an object, its shape, size, speed, and direction [7], [34]. Having such a forward hydrodynamic model does not directly allow for decoding the velocity profile, i.e. determine the mentioned properties of an object. This defines the inverse problem: reconstructing an object's properties from measured velocity profiles.

One aspect of these spatial velocity profiles, notably spatial broadening [7], [11], has been shown to be beneficial with respect to localizing vibrating spheres, a subset of the full inverse problem. The spatial broadening property entails that the generic shape of a velocity profile is consistent when an object is moving in a certain direction. This generic velocity profile shape, or signature, is only scaled and translated by changes in the relative position, size and speed of the object. This property has been successfully used in template matching methods for source localization [7], [35]. More directly, zero crossings, maxima, and minima of the velocity profile can also be directly used to produce estimates for the distance (y-coordinate) and the lateral position (x-coordinate) [10]. In [30], [36], [37], it is shown that an ALL comprising of

2D-sensitive sensors makes localization more robust. Compared to the 1D sensitivity of fish and most other ALL implementations, this extension provides more and complementary spatial reference points [30], [36], which may also be helpful for shape recognition.

2) HYDRODYNAMIC SIGNATURES

Some ALL systems focus on estimating fluid flow parameters such as the flow speed, direction and vorticity, either in natural environments [38] or in confined flow tanks [39]. In the latter case, an object is usually placed upstream with respect to a sensor array [18]. At constant flow rates, this object can shed Kármán vortex streets that encode object characteristics in the spatial and periodic properties of the vortices [19], [21], [22], [34].

The alternative approach, that we use in the current study for object shape classification, is to move the object rather than the medium. This allows measuring the hydrodynamic effects of the object shape. In the hydrodynamic near-field, the shape of the object affects the pressure gradient and thus also the principal shape of the measured velocity profile. This measured velocity profile shape can be modeled via a process called conformal mapping [34], [37], although these velocity profiles are expected to be less distinctive at distances further from the array [20] and in the limit resemble the signature of a sphere. So in order to preserve details and identify objects based on their measured velocity profiles, they need to be measured up close. This allows constructing what we define as hydrodynamic signatures of each object in detail.

3) STATE-OF-THE-ART FLOW-BASED SHAPE CLASSIFICATION

There are several works that have demonstrated object shape classification in simulation or small-scale experiments.

Recently [37], neural networks have been used in a simulation study to determine the shape parameters of a foil-shape object. Using a grid of sensors and a conformal mapping potential flow fluid model, the two perpendicular fluid flow components (x and y) were simulated, as well as the absolute fluid speed magnitude and the dynamic pressure at each of the sensors. The shape of the object was reconstructed to a high degree via a triplet of neural-network estimated shape parameters. The dynamic pressure magnitude was shown to the best choice of the four considered types of input to the neural network.

In [18], a pressure sensor array is used to determine the shape parameters of a cylinder and cube object in a fluid flow channel. From the reported statistics, we infer an F1-score of 97.1% for this two-class problem. In a similar study [19], two object shapes (semi-circle and rectangle) of three sizes were analyzed via their wake. Visualizing the dominant wavelength of the vortex streets versus its magnitude reveals that measurements on objects of similar size cluster together. Discerning the shape would however require additional features.

In [23], a stationary array was used in quiescent water to infer the shape parameters of an elliptical cylinder. The objects were moved past the array at a constant speed of 50 cm/s at a constant distance between 0.5 and 1 cm. Between runs, the angle relative to the array and radii of the elliptical cylinder were varied. By employing a particle filtering method, they were partly able to reconstruct these two types of shape parameters.

Compared to these related works, we increase the distance and size of the objects, the number of different object shapes, and the scale of the array. In addition, we measure the flow field in two dimensions; novel to the application of shape recognition via hydrodynamic imaging.

4) TEMPORAL VELOCITY PROFILE SENSING

With small scale or high density sensor arrays, neighboring sensors often show correlation which is harnessed in the signal processing pipeline. This assumes that the sampled velocity profile has an adequate spatial resolution to be matched to a spatial hydrodynamic signature of a known object. For large-scale arrays and applications, this correlation assumption breaks down.

Instead of treating the measurements from the sensor array as under-sampled spatial velocity profiles, we switch to the time-domain and focus on temporal velocity profiles, in our case measured from an object that moves with constant speed past the sensor array. While additional sensors would likely improve the quality of the measurements, one beneficial consequence of sparse spatial sampling is that we have true independent measurements of an event, where hydrodynamic noise can thus be factored out more easily. In fact; we have eight independent measurements across the array.

B. ELM NEURAL NETWORK

For classification, we employ a light-weight regression artificial neural network, the Extreme Learning Machine (ELM) [26], which has been shown to be fast and straightforward to optimize for ALL localization [11], [12], [40] without overfitting.

While other, more advanced neural networks might be more powerful, determining optimal network configurations and preventing overfitting for these more complex neural network architectures is not trivial. The ELM has only one network parameter to choose: its hidden layer size, i.e. the amount of hidden neurons, which makes it straightforward to tune to avoid overfitting.

The second benefit of the ELM neural network is that it can be trained in a fraction of the time needed for an MLP (multilayer perceptron) or deep neural network. The ELM has only one hidden layer, with weighted connections to the input layer and output layer. The fast training of this network is a result of only the output weights being trained; the input weights remain fixed after initialization [26]. This allows the network to be trained in a single learning step, while still providing a non-linear transformation of the input space.

III. METHODS

We measure temporal velocity profiles using a large-scale ALL in response to differently shaped objects attached to a

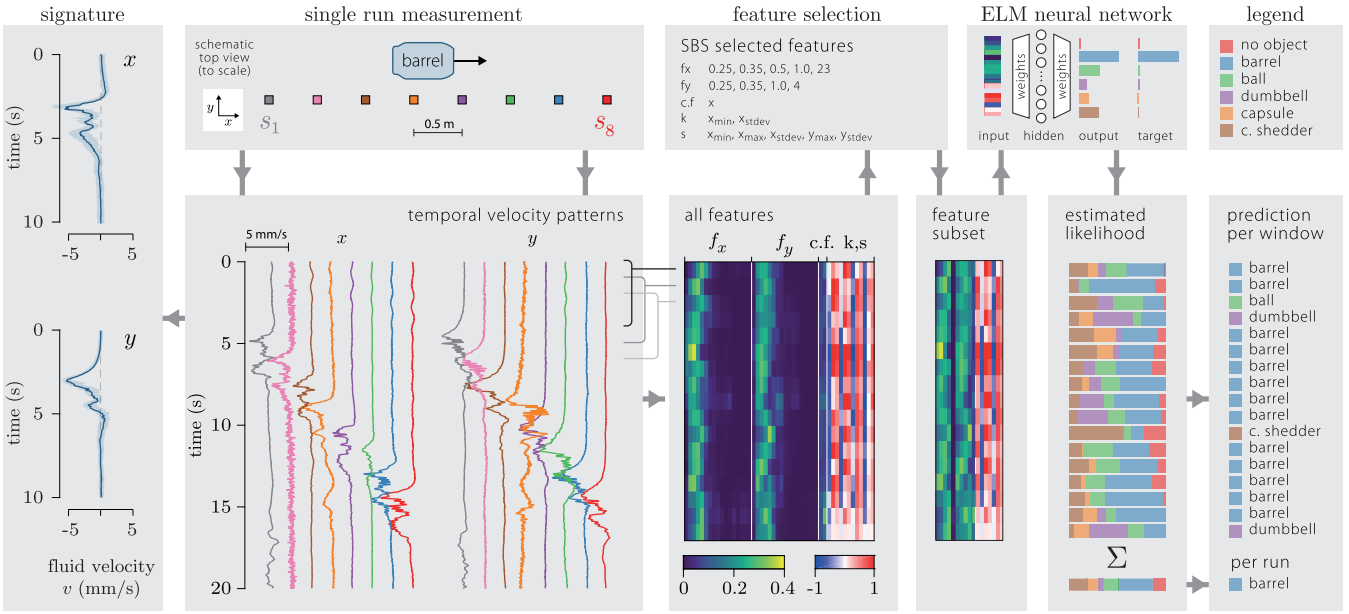

FIGURE 2. Signature and classification pipeline overview for a forward moving barrel. Arrows between the blocks show the flow of signal processing. The single run is measured by the ALL and yields temporal velocity profiles for both sensing dimensions. These are time-shifted and aligned to produce a signature in the leftmost pane. For classification, we calculate features of the flow based on overlapping 4 second windows. Here, $f_{x,y}$ denote the frequency bands, `c.f.` the centroid frequencies, and `k` and `s` the kurtosis and skewness features respectively. We use feature selection methods to create feature subsets, the result of the SBS method is shown here. The ELM is trained to match each feature (input) vector to a target vector indicating the correct class. When the trained ELM is presented with a new feature vector as input, the network outputs likelihoods for each class. We obtain shape predictions by selecting the maximal likelihood for each time window, or for the whole run.

towing platform. These profiles are used to visualize hydrodynamic signatures, as well as input for the classification pipeline. In this pipeline (see Fig. 2), we first calculate features, then determine feature subsets, and finally evaluate the classification performance using these subsets via the ELM neural network. These processes are described in detail in the following subsections.

A. SETUP

Here, we first describe the sensor array, followed by the towed objects, and finally the experimental procedure.

1) ALL SENSOR ARRAY

To demonstrate object shape classification, we installed eight 2D-sensitive all-optical flow sensors [12], [24] each 0.5 m apart at the short side of a 18 m × 25 m swimming pool, see also Fig. 4. The 3.5 m array was placed 0.5 m from the wall and centered with respect to the 18 m wide pool area. The sensor array has an adjustable submerged depth, which defaults at 1.4 m. The array is adjusted in depth for every submerged object to be level with its center, as to accurately measure the hydrodynamic signature of each object.

Each all-optical flow sensor is constructed using a fiber structure comprising four optical fibers. Since they are all-optical, they do not directly require electricity to function, but rather operate on optical signals using fiber Bragg gratings (FBGs) [24]. The white spherical element at the tip of each sensor (Fig. 4, right) picks up fluid forces and deflects the fiber structure. This deflection is measured using the FBGs inside the fiber structure which can be read out at a distance via optical cables.

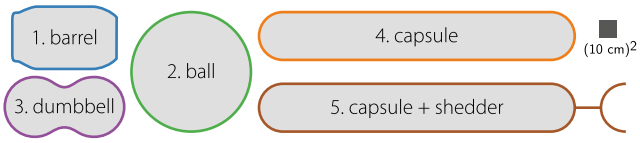

FIGURE 3. Schematic side view (to scale) of the objects used.

The 32 optical signals from 8 sensors were simultaneously measured using a Hyperion si-225 optical interrogator (Micron Optics, USA) at 5 kHz, and downsampled to 200 Hz. We report and graph the sensed local fluid velocity for each of these sensors in the x-dimension (parallel to the array) and y-dimension (orthogonal to the array).

2) TOWED OBJECTS

We measure the temporal velocity profiles resulting from different objects towed past the ALL. In addition, for a null measurement baseline, we measured velocity profiles with only the towing platform as the 'no object' case.

The towing platform is constructed from two floating pontoons 70 cm apart, each 16 cm in diameter and 150 cm in length, together providing 60 L of buoyant volume. An electric winch was used to tow the platform and attached object with a constant speed, specified to maximally tow 200 kg at 0.3 m/s.

Five differently shaped objects were chosen for the experiments (Fig. 3). A dumbbell and barrel shape were chosen since they displace roughly equal amounts of water, yet have a different shape. A ball was chosen as it matches the spherical shape used in theory and most ALL characterization

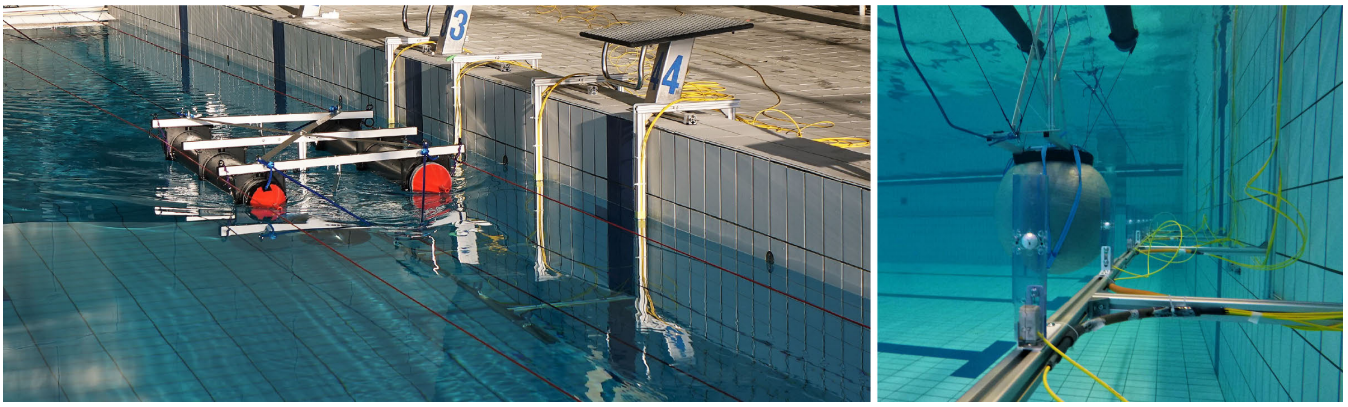

FIGURE 4. Overview of the deployed setup. The left picture shows the towing platform in motion near the center of sensor array with two (red) guiding wires and a partly submerged (blue) towing cable. The right picture shows the sensor array with the white spherical elements at a depth of 140 cm with respect to the water level, as well as the ball object attached to the towing platform.

TABLE 1. List of run distances, object sizes, and estimated Reynolds numbers (Re) based on their displaced volume.

Object	Distances (m)	Size (cm)	Re
Baseline	0.4	-	-
Barrel	0.4, 0.57	⌀37×63	15×10^4
Ball	0.6, 0.84	⌀70	21×10^4
Dumbbell	0.4, 0.57	⌀35×70	14×10^4
Capsule	0.4	⌀28×185	18×10^4
C. + shedder	0.4	⌀28×(185+30)	19×10^4

efforts [4]; moreover it displaces a lot of water, producing a strong signal. Finally, two variants of a capsule shape were deployed. First a slender capsule, and secondly the same capsule with an attached inverted bowl, which acts as a wake-shedder. All shapes were mounted a meter under the towing platform and were made neutrally buoyant by filling them with water.

3) EXPERIMENTAL PROCEDURE

Throughout the rest of this paper, we consider a 'run' to be a completed motion of the towing platform, towing an object past the array once. For each object, we measure three runs from left to right (forwards) and three runs from right to left (backwards) at each distance as indicated in Table 1. With an average run duration of 50 seconds to cross the pool, we selected 20 seconds from each run in which the towing platform passes by the array.

The inviscid hydrodynamic model [31], [32] as described for 2D sensing in [11] allows determining an upper bound on the detection distance, depending on the stimulus in question. For a conservative estimate, we consider a 0.2 m radius object moving parallel to the array at 0.3 m/s, and a noise-level equivalent sensor threshold of 5 mm/s for DC signals [24]. For a signal to noise ratio of one, such an object can be detected with the center at 0.62 m distance from the sensor array. To preserve more details for the hydrodynamic signature of these objects, we vary the object center distances around a closer value, see also Table 1.

B. SIGNAL PROCESSING

We process the data in two ways. First we align and combine all measured temporal velocity profiles to construct and visualize temporal hydrodynamic signatures. In contrast, for the object classification framework, we use the eight unaligned and unfiltered 2D velocity profiles.

1) ALIGNING FOR HYDRODYNAMIC SIGNATURES

To construct the signatures, we combine 24 temporal velocity profiles from 8 sensors and 3 repeated runs for each object, distance, direction, and sensing dimension. We first low-pass filter the signals (Butterworth, 3 Hz) and use correlation to find the delays between sensors. Then we time-shift the original, unfiltered velocity profiles to correct for these delays. Finally, we show the median signature, as well as the 25 to 75 percentiles as an indication of the consistency of the measured signatures.

2) FEATURE EXTRACTION

In our classification pipeline, we make use of a moving window approach. Here we use a time window of 4 seconds and a stride of 1 second to calculate features of the measured fluid flow. This length was chosen since characteristic wakes and hydrodynamic stimuli are in the <5 Hz range and are well-captured in this time-scale. For all features, we calculate the feature values for the x and y sensing dimension separately.

We consider two types of features, spectrum based and distribution based. An overview can be found in Table 2; we briefly describe them here.

We include the DC offset (mm/s) within a time window as the first spectral feature. The energy in different frequency bands proved to be an important feature for classifying river flow conditions [38]. We therefore select 16 frequency bands from the ANSI half-octave band definition [41], using 0.25 Hz as a reference frequency. This yields mid band frequencies from 0.25 Hz to 45 Hz on a logarithmic scale. We used 16 2nd-order Butterworth filters to filter each windowed signal and determine the frequency band energies.

TABLE 2. List of used features which are calculated separately for the *x* and *y* dimensions, combining to 48 features in total.

Feature	Description	Amount
DC	0 Hz	1
midband frequency of f-bands	{.25, .35, .50, .70, 1.0, 1.4, 2.0, 2.8 4.0, 5.7, 8.0, 11, 16, 23, 32, 45} Hz	16
c.f.	-	1
kurtosis	{min ($\vee$), max ($\wedge$), σ}	3
skewness	{min ($\vee$), max ($\wedge$), σ}	3

The final spectral feature we use is the centroid frequency of the average spectrum in a time window. For all spectral features, we average the feature value over the sensors to a single feature value per time window.

The second feature type, i.e. kurtosis and skewness, describes the distribution of data in the time windows, as used before in fluid flow classification [38] and sensor placement optimization [42]. Instead of averaging over the eight sensors, we take the minimal value, maximal value and standard deviation of both as separate features.

3) FEATURE NORMALIZATION

Before selecting the features and feeding them to the classifier (section III-C1), the feature vectors are normalized to have their feature values in a more favorable range for the classifier.

For the frequency bands, we take the spectral power for each band and divide this by the bandwidth of their respective band. We further normalize the frequency band features for each time window by dividing their values by the total power for that time window. This process re-scales all values to the range from 0 to 1 and makes their sum 1. These features are now less dependent on the actual flow speed, but rather encode relative contributions, making them more versatile with respect to object size and speed.

For kurtosis and skewness we selected the minima, maxima and standard deviation over the eight sensors. We normalize these by applying the hyperbolic tangent, scaling these distribution feature values between -1 and 1.

4) FEATURE SELECTION

Some of the selected features might still be irrelevant or encode duplicate information. To aid the classifier in preventing overfitting, we search for subsets of features that are relevant for discerning the objects. We consider three types of methods for selecting feature subsets [25]: manual selection, filter methods, and wrapper methods. By showing the performance of these subsets, we can infer which features are more informative for classifying objects.

We first demonstrate the performance of the classifier with handpicked subsets, in our case: all 48 features, frequency bands, only x frequency bands, only y frequency bands, and finally the centroid frequencies combined with kurtosis and skewness.

Filter methods usually select features based on their similarity or relevance, defined through correlation. These methods are unsupervised, in the sense that they do not make use of the information about which measurement belongs to which class. For the filter methods, we use Lasso, reliefF [43], FCBF [44], mRMR [45], and cmRMR [46]. Two of these methods require the features to be discretized in bins, others work on continuous values. When binning is required, we define 16 sequential bins per feature, all equally populated. This is achieved via determining its distribution, and splitting the distribution on the median of the feature values. We repeat this step four times to obtain the bin boundaries. We only use binning for selecting the features; for the final classification, the original continuous feature values are used.

Wrapper methods usually test different combinations of features and monitor the classification performance; a supervised method. For the wrapper method, we use a $\ell 1$-svm (support vector machine) as a linear classifier, as suggested by [47] to provide an unbiased wrapper. Were one to use the final classifier, one could tailor the feature selection process too much and thus cause overfitting, reducing generalization. With this wrapper, we perform both a sequential backward search (SBS), eliminating one feature at a time; and a sequential forward search (SFS), adding one feature at a time, to produce a subset of features.

C. CLASSIFICATION

We measure three repeated runs per object, distance, and motion direction, resulting in a total of 54 measured runs. We employ 3-fold stratified cross validation: each fold contains a complete set of 18 runs. Two folds are used for training and tuning the neural network to avoid overfitting, whereas the remaining fold is used for testing.

1) ELM NEURAL NETWORK TRAINING

As described in the Background section, the ELM has a single tunable parameter: the size of the hidden layer. To determine the optimal size before each training phase, we employ a nested validation scheme. During this tuning-phase, the data in the training set is further randomly divided into 5 sub-folds for nested cross validation, four of which are used for training the network. We use the remaining sub-fold in this tuning phase for monitoring the performance while varying the network hidden layer size.

For the training phase, we then take the five found optimal sizes during the tuning phase and use the average of these sizes to retrain the network on the whole training set. When the network is trained, we test it on the remaining fold for the exploitation phase.

As for the network structure, the trained network's input size is determined by the number of selected features. Its hidden layer size is determined during the tuning phase, which was usually between 40 and 50 nodes. The output is always a 6-element vector: a likelihood score for each class, also known as one-hot encoding.

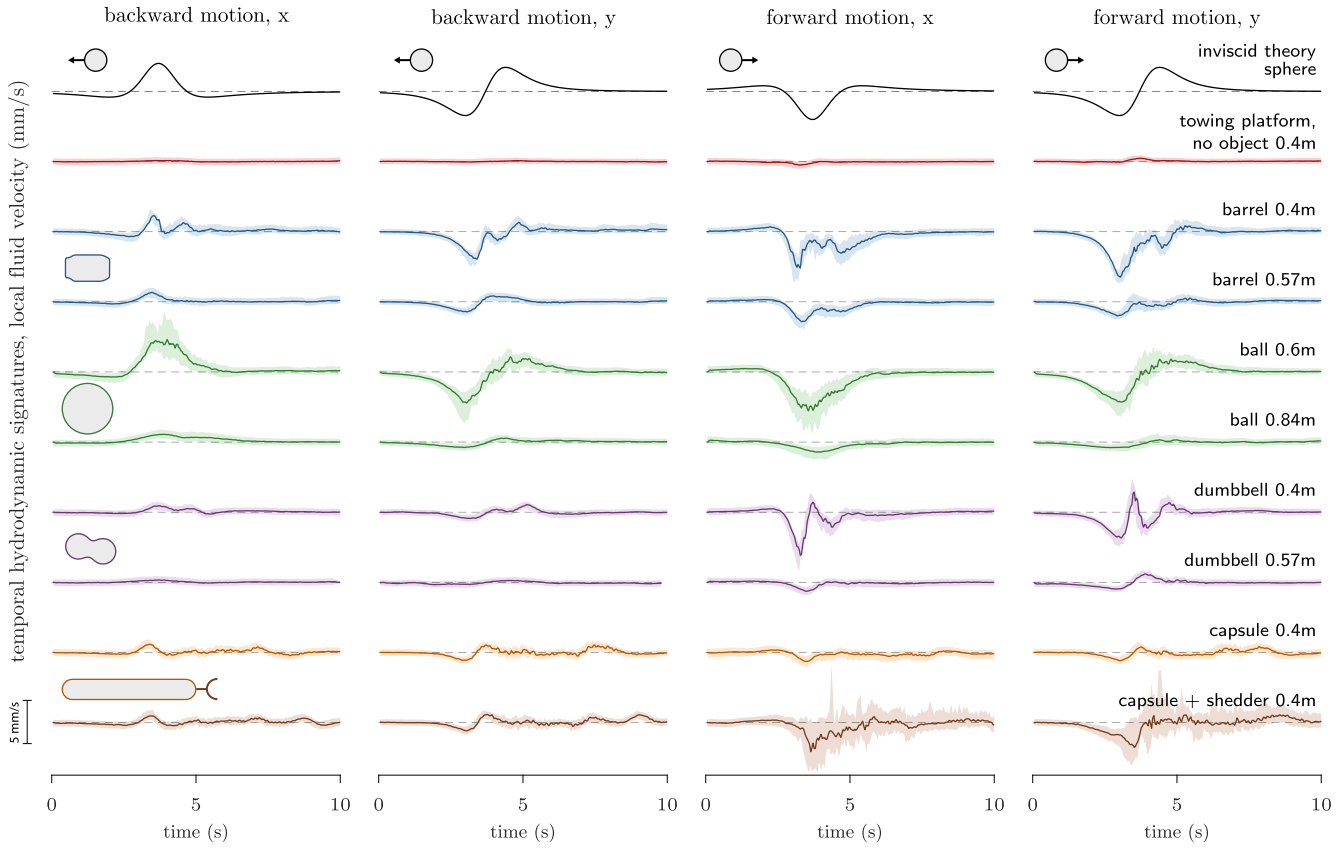

FIGURE 5. Constructed temporal hydrodynamic signatures of each object and distance (right pane). The object shapes are indicated to relative scale in the left pane. The signatures are visualized with the median response and the 25th to 75th percentiles (fills), which are calculated from 24 time-aligned temporal velocity profiles.

2) CLASSIFIER PERFORMANCE

Because there is a separate repetition of each type of run in each fold, we can assert the performance consistency. Defined as usual, we report the classification performance as the precision, recall, and average F1-score. We have two modes of scoring the classifier performance.

First, we report the classification performance 'per window'. Here we simply take the maximal value of the output vector as the predicted object for each time window.

The second mode of reporting, 'per run', first averages the output vectors of the 17 time windows in a run, then taking the maximal value as a single prediction for the whole run. With this approach, our classification pipeline has the attractive property that it can use mounting evidence and can thus be used on signals with arbitrary length.

IV. RESULTS

We first discuss the constructed temporal hydrodynamic signatures and reflect on their apparent relation to the object shapes. Secondly, we report the shape classification performance for different feature subsets, which indicate what types of features are suited for discriminating between object shapes.

A. TEMPORAL HYDRODYNAMIC SIGNATURES

We measured the velocity profiles resulting from moving five different objects past our array. Including the empty towing platform as a baseline, we have a total of six classes. Fig. 5 shows constructed signatures, as well as the theoretical signature for a sphere in inviscid fluid [31] as described for 2D in [11]. We show the x- and y-component separately and make a distinction between the direction of motion and the measured distances.

We measured some objects at two distances, to show the effect of distance on the measured hydrodynamic signatures. As expected from conformal mapping theory, the signatures at the further distances are indeed lower in magnitude and lose detail compared to their respective nearby version. When further away, these object signatures start to resemble the signature of a sphere, losing higher order components to the noise. These objects are still detectable, but not immediately identifiable with respect to their shape.

The complementary effect of 2D sensing is also visible in these temporal hydrodynamic signatures. When e.g. the x sensing value reaches a zero-crossing, the y dimension often shows a local extrema or other non-zero value and vice versa. This increases the signal to noise ratio of the combined measurement at these informative points.

FIGURE 6. Bar plots indicating the 3-fold classification performance on the train and test sets for several feature subsets. Hand picked sets are grouped left, filter method found sets are grouped right, and the two wrappers are displayed at the bottom. The 'test per run' score and standard deviation are displayed next to each bar plot, other scores are included in Table 3. The best performing hand picked and algorithmically determined subset are indicated in bold. Methods with a ∗ make use of a discretised version of the features during the selection process.

Overall, the signatures resulting from forward and backward motion are quite similar in shape. The notable exceptions here are the dumbbell object and the capsule shedder object. The asymmetry in the dumbbell's signatures is likely caused by the object being mounted at a slight angle during experiments. For the capsule shedder object, the forward motion clearly causes turbulent wake shedding to occur.

Furthermore, several identifiable relations emerge between the object shapes and their measured hydrodynamic signatures (Fig. 5). First, the ball-shaped object produces a signature similar to that of a modeled sphere in inviscid flow. A second relation is visible between the barrel and dumbbell shape. Both signatures are similar in that they show two minima and maxima, where the dumbbell has a sharper signature, especially during forward motion. Thirdly, the effect of the wake shedder on the capsule is clearly visible during the forward motion, producing turbulent behavior. In addition, during backward motion, the signatures of these two shapes are similar and match quite well. They differ in a subtle extra bump at the 9 second mark, coinciding with the addition of the inverted bowl.

These observations suggest that the object shape is reflected in their temporal hydrodynamic signature.

B. SHAPE CLASSIFICATION

For shape identification, we only considered the original unaligned sensor measurements and the object shape, ignoring the distance and relative location of the object.

1) CLASSIFICATION PIPELINE

One example of the process of a classified run is visualized in Fig. 2. From the visualized feature vector, we observe that the energy bands in the lower part of the spectrum contain the most energy, as expected. Furthermore, the spectral feature values remain similar throughout the subsequent time windows of the run.

For for 4 out of 17 time windows in this example run (Fig. 2), the barrel object did not receive the highest neural network output. If we however aggregate this output over the whole run, the barrel shape clearly has the highest predicted likelihood and is correctly predicted as a barrel using the 'per run' approach.

2) FEATURE IMPORTANCE

Fig. 6 provides a visual indication of the classification performance for each subset, as well as the exact achieved classification test score. A full list of chosen features for each feature subset and exact scores can be found as Table 3.

The influence of each feature subset on overfitting is also visible in Fig. 6 via the discrepancy between 'train' and 'test' performance. A large discrepancy indicates overfitting and loss of generalization. Quite notable is the second feature subset 'all f-bands', where seemingly no overfitting takes place.

The performance with the handpicked feature subsets, as indicated in Fig. 6, show that the spectrum features (all f-bands) yield higher performance than the other features combined (c.f. & k,s). In addition, the velocity profile orthogonal to the array, y, seems more informative for discerning the objects, compared to the traditional parallel component x, as sensed by the fish lateral line and as used in most 1D ALL applications [4].

For the algorithmically generated subsets, the quantitative contribution of each feature to the final classification performance is hard to estimate, given that one also has to consider the interaction and overlap between the features.

TABLE 3. List of features subsets, the number of features, and the 3-fold performance per single time window (w) and per whole run (r). The F1-score of the best performing hand picked and algorithmically determined subset is indicated in bold. The feature selection methods with a $*$ make use of a discretised version of the features during selection. The shorthand symbols (minima, maxima, standard deviation) for kurtosis and skewness coincide with those of Table 2.

method for feature subset	f_x (Hz)	f_y (Hz)	c.f.	k+s	#	train (w) F1 $\pm\,\sigma$ (%)	test (w) F1 $\pm\,\sigma$ (%)	train (r) F1 $\pm\,\sigma$ (%)	test (r) F1 $\pm\,\sigma$ (%)
all	DC to 45	DC to 45	x, y	all	48	74.8 ± 0.59	64.8 ± 1.23	100 ± 0	96.3 ± 1.86
all f-bands	DC to 45	DC to 45	-	-	34	69.8 ± 1.77	61.9 ± 1.38	98.6 ± 0.77	$\mathbf{98.6 \pm 0.99}$
x f-bands	DC to 45	-	-	-	17	58.4 ± 0.75	51.7 ± 1.06	80.0 ± 1.93	75.7 ± 3.32
y f-bands	-	DC to 45	-	-	17	58.0 ± 1.38	51.3 ± 1.02	95.0 ± 1.72	87.6 ± 3.41
c.f. & k+s	-	-	x, y	all	14	60.0 ± 0.56	52.6 ± 1.42	95.2 ± 0.88	90.8 ± 1.88
Lasso	0.5, 1.4, 2.0, 2.8, 4.0, 5.7	0, 0.7, 1, 2.0, 11	-	$k_{x\vee}$, $k_{x\sigma}$, $s_{x\vee}$, $s_{x\sigma}$, $s_{y\wedge}$	16	63.5 ± 0.82	55.0 ± 0.69	96.4 ± 0.51	85.7 ± 0.98
reliefF [43]	0.25	0.25, 0.35	-	all $\neg$ ($k_{x\sigma}$, $s_{x\sigma}$)	13	59.5 ± 1.45	51.8 ± 1.18	90.6 ± 2.85	78.6 ± 5.95
mRMR* [45]	-	0, 0.25, 4.0, 8.0	x	$s_{y\wedge}$, $s_{y\sigma}$	7	66.4 ± 0.88	62.6 ± 1.07	93.8 ± 1.05	88.3 ± 1.35
FCBF* [44]	1.0	0.25, 4.0	x	$k_{x\vee}$, $k_{y\vee}$, $s_{x\vee}$, $s_{y\wedge}$	8	66.7 ± 0.87	60.9 ± 1.08	98.0 ± 0.72	94.6 ± 1.47
cmRMR [46]	0, 1.0, 2.8, 4.0, 2.0, 16, 45	0, 0.7, 1.4, 11, 45	y	$s_{x\vee}$, $s_{y\sigma}$	15	65.7 ± 1.34	0.4 ± 0.65	95.2 ± 1.83	91.0 ± 2.55
SBS [47]	0.25, 0.35, 0.5, 1.0, 23	0.25, 0.35, 1.0, 4.0	x	$k_{x\vee}$, $k_{x\sigma}$, $s_{x\vee}$, $s_{x\wedge}$, $s_{x\sigma}$, $s_{y\wedge}$, $s_{y\sigma}$	17	71.7 ± 1.04	64.1 ± 0.79	100 ± 0	$\mathbf{95.8 \pm 0.61}$
SFS [47]	0, 0.25, 0.50, 23	0.35, 1.0, 2.0, 5.7	-	$k_{x\vee}$, $k_{y\vee}$, $s_{x\sigma}$, $s_{y\vee}$, $s_{y\wedge}$, $s_{y\sigma}$	14	69.8 ± 0.59	62.6 ± 0.50	97.1 ± 0.91	94.6 ± 3.38

Some feature selection algorithms provided subsets of features which led to poor performance consistency, such as the Lasso and reliefF methods. These methods provided subsets that perform well on the training set, but show a considerable drop in performance for the test set, which indicates poor generalization.

For our experiments, the hand picked 'all f-bands' subset (98.6%) and the ℓ_1-SBS (sequential backward search, 95.8%) algorithm provide the best feature subsets; both show a high testing performance as well as a low discrepancy between training and testing performance, indicating good generalization.

V. DISCUSSION

The sensor array of all-optical 2D-sensitive flow sensors allowed us to capture a profile of the hydrodynamic environment at eight measurement points in 2D. We have demonstrated that the sparse large-scale 2D-sensitive ALL enables hydrodynamic shape identification with high precision.

By carefully tuning the ELM neural network classifier via nested cross-validation and employing feature selection, we kept overfitting to a minimum, as signified by the small discrepancies between train and test scores. This indicates that the methods and features used may generalize well to different sets of measurements.

A. HYDRODYNAMIC SIGNATURES

As mentioned in the introduction and background sections, our sparse large-scale ALL has the benefit of independently sampling eight temporal velocity profiles. We realigned these profiles to construct temporal hydrodynamic signatures (Fig. 5), which can be used for further analysis.

The temporal hydrodynamic signatures show that the object shape can be reflected in its signature. From Fig. 5, we observe that similarly shaped objects also show visual similarities in their measured hydrodynamic signatures, such as between the barrel and dumbbell shape and between the two capsule objects.

The hydrodynamic signature of the towed ball (Re $\approx$ 210000) closely resembles that of a modeled sphere in an inviscid fluid. This is remarkable, since the towed object generates a flow regime where vortices occur which would normally introduce chaos. While the vortices are still present in the data, as evidenced by the variation (percentiles, fills) in Fig. 5, combining these measurements produced a smooth median signature.

Although the temporal hydrodynamic signatures are useful for visual inspection, the method of time-shifting signals is not well suited for classification, as it requires a constant speed and completed runs. To introduce more flexibility with respect to these constraints, we used a different processing method for the classification process.

B. FEATURE AND WINDOW BASED CLASSIFICATION

The approach of selecting feature subsets allowed optimizing the classifier, whereas the aggregating window approach made the method location-invariant and flexible towards measurement duration.

From the full list of selected features for each subset (Table 3), some trends emerge. In general, the lower frequency bands ($<$ 4 Hz) are selected more often than the higher frequency bands. There is no significant difference between the amount of selected x versus y dimension features. The distribution features, kurtosis and skewness, are

furthermore often chosen as informative during the feature selection process. This might be due to the low frequency shedding characteristics, which are less visible in the average hydrodynamic signatures (Fig. 5) given their quasi-random nature, but are more visible in individual temporal velocity profiles (Fig. 2).

The best performing feature subsets were the handpicked set containing both x and y frequency bands (98.6 ± 0.99 %), and the set of 17 features found via sequential backward search [47] (SBS, 95.8 ± 0.61 %). Another feature subset that stands out is the one found through FCBF [44] (94.6 ± 1.47 %), which only selects 8 features to achieve this classification performance. This illustrates that selecting informative features can be equally or more important compared to optimizing the classifier.

There are however some situations in which the classifier makes mistakes based on information from single four second time windows. From Fig. 7A it is clear that there is some confusion between the barrel, ball, and dumbbell shapes. This has likely to do with that these objects were measured from two distances, where objects further away produce more generic velocity profiles, as evidenced by the constructed hydrodynamic signatures and as expected from literature [20], [34]. For the capsule objects that were measured at a single distance, we also see some confusion, albeit less.

When we apply a majority vote per run, only few mistakes remain as indicated in Fig. 7B, regardless of the object's distance. The classification performance thus improves from mounting evidence, effectively averaging the neural network prediction vector over sequential time windows in a run. The only mistakes that remain with a 'per run' classifier (Fig. 7B) is that a dumbbell is mistaken for a barrel in one instance, and the capsule with a shedder is mistaken for a capsule without a shedder in another. Given the similarities between these shapes, these wrong predictions can still be considered informative.

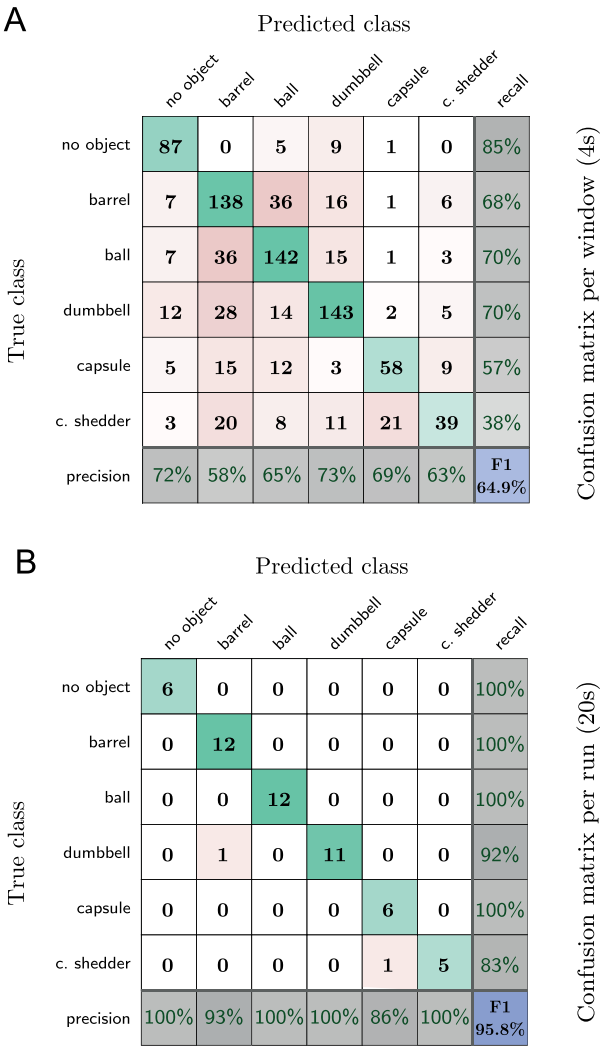

FIGURE 7. Confusion matrices for classification using the SBS feature set. (A) shows the predictions per window, whereas (B) shows the predictions after taking the per run approach.

C. 2D SENSING

The hydrodynamic signatures demonstrate the complementary advantage of 2D sensing; at zero crossings or low magnitudes in one sensing dimension, the other dimension often shows a local extrema when an object is nearby. This complementary effect is also visible from the final performance, where the best performance was achieved when combining both sensing dimensions.

Via the hand-picked feature subsets (Fig. 6), we find that the novel y component is more informative for discerning between objects. Combining the two measuring dimensions improves classification even further, which also holds for object localization [11], [36].

D. FUTURE RESEARCH

Since the found feature subsets are tailored to discriminate between the considered object shapes, some features that were not chosen may still prove instrumental in discerning other shapes that may be added to a future data set or experiment.

For expansion on this research, and to further attest the utility and robustness of shape identification via hydrodynamic imaging, it would be interesting to combine the two current strategies for measuring the hydrodynamic effects of shapes. Both a stationary object placed upstream in a constant speed fluid flow, as well as an object moving with constant speed in still water, selectively consider one aspect of fluid flow sensing. Both the properties in Kármán vortex streets from literature and temporal hydrodynamic signatures presented here, contain information about the object. As a first step, one could include an external flow source during measurements on moving objects, making for more realistic conditions.

Although the objects were moved at a single constant speed and at a constant distance, our classification pipeline does allow for deviations from this scheme. Due to the nature of the near-field hydrodynamic imaging, the objects should be close to the array, but not necessarily at a constant distance or for a fixed duration. Our processing pipeline ensures that the

signals can be of arbitrary length and from any of the sensors from the array, and still produce a correct classification at a 4× chance rate on a single four second time window, as indicated by our 'per time window' scores (Fig. 7A). Taking a four second buffer into account, this pipeline could therefore be potentially be implemented online to provide online shape classification.

VI. CONCLUSION

By scaling up the artificial lateral line (ALL), we have demonstrated that near-field hydrodynamic imaging can be used on a supra-biological scale. When deployed as a static array, the system can be used for near-field object tracking and, as demonstrated here, object identification.

Furthermore we have not only shown that 2D sensing is beneficial with respect to shape identification, but also that the novel orthogonal velocity component bears more information that the traditional parallel velocity component, as used by fish and most 1D ALL applications.

This demonstrates that this 2D passive technology is especially suited to identify a moving object via its hydrodynamic interactions, and suggests that it can be used for object detection, object identification, and collision avoidance in dark and murky environments.

ACKNOWLEDGMENT

The authors would like to thank A. Dzipalski, S. Warmelink, S. Römer, and the Hanze Institute of Sport Studies for their contribution in facilitating this study.

DATA AVAILABILITY

The data sets generated and analyzed during the current study are available from the Dataverse repository [48] or by contacting the corresponding authors.

REFERENCES

[1] S. Dijkgraaf, "The functioning and significance of the lateral-line organs," *Biol. Rev.*, vol. 38, no. 1, pp. 51–105, Feb. 1963.

[2] H. Bleckmann, J. Mogdans, J. Engelmann, S. Kröther, and W. Hanke, "Wie fische wasser Fühlen: Das seitenliniensystem," *Biologie Unserer Zeit*, vol. 34, no. 6, pp. 358–365, 2004.

[3] S. M. van Netten and M. J. McHenry, "The biophysics of the fish lateral line," in *The Lateral Line System*, vol. 48, S. Coombs, H. Bleckmann, R. R. Fay, and A. N. Popper, Eds. New York, NY, USA: Springer, 2013, pp. 99–119.

[4] G. Liu, A. Wang, X. Wang, and P. Liu, "A review of artificial lateral line in sensor fabrication and bionic applications for robot fish," *Appl. Bionics Biomech.*, vol. 2016, pp. 1–15, 2016.

[5] Y. Yang, J. Chen, J. Engel, S. Pandya, N. Chen, C. Tucker, S. Coombs, D. L. Jones, and C. Liu, "Distant touch hydrodynamic imaging with an artificial lateral line," *Proc. Nat. Acad. Sci. USA*, vol. 103, no. 50, pp. 18891–18895, Dec. 2006.

[6] A. T. Abdulsadda and X. Tan, "An artificial lateral line system using IPMC sensor arrays," *Int. J. Smart Nano Mater.*, vol. 3, no. 3, pp. 226–242, Sep. 2012.

[7] B. Curcic-Blake, "Source location encoding in the fish lateral line canal," *J. Experim. Biol.*, vol. 209, no. 8, pp. 1548–1559, Apr. 2006.

[8] M. Asadnia, A. G. P. Kottapalli, J. Miao, M. E. Warkiani, and M. S. Triantafyllou, "Artificial fish skin of self-powered micro-electromechanical systems hair cells for sensing hydrodynamic flow phenomena," *J. Roy. Soc. Interface*, vol. 12, no. 111, Oct. 2015, Art. no. 20150322.

[9] X. Zheng, Y. Zhang, M. Ji, Y. Liu, X. Lin, J. Qiu, and G. Liu, "Underwater positioning based on an artificial lateral line and a generalized regression neural network," *J. Bionic Eng.*, vol. 15, no. 5, pp. 883–893, Sep. 2018.

[10] J.-M. P. Franosch, A. B. Sichert, M. D. Suttner, and J. L. V. Hemmen, "Estimating position and velocity of a submerged moving object by the clawed frog Xenopus and by fish—A cybernetic approach," *Biological*, vol. 93, no. 4, pp. 231–238, 2005.

[11] B. J. Wolf and S. M. van Netten, "Hydrodynamic imaging using an all-optical 2D artificial lateral line," in *Proc. IEEE Sensors Appl. Symp. (SAS)*, Sophia Antipolis, France, 2019, pp. 1–6.

[12] B. J. Wolf, S. Warmelink, and S. M. Van Netten, "Recurrent neural networks for hydrodynamic imaging using a 2D-sensitive artificial lateral line," *Bioinspiration Biomimetics*, vol. 14, no. 5, Jun. 2019, Art. no. 055001.

[13] L. Devries, F. D. Lagor, H. Lei, X. Tan, and D. A. Paley, "Distributed flow estimation and closed-loop control of an underwater vehicle with a multi-modal artificial lateral line," *Bioinspiration Biomimetics*, vol. 10, no. 2, Mar. 2015, Art. no. 025002.

[14] M. Krieg, K. Nelson, and K. Mohseni, "Distributed sensing for fluid disturbance compensation and motion control of intelligent robots," *Nature Mach. Intell.*, vol. 1, no. 5, pp. 216–224, May 2019.

[15] A. N. Vollmayr, S. Sosnowski, S. Hirche, and L. J. van Hemmen, "Snookie: An autonomous underwater vehicle with artificial lateral line system," in *Flow Sensing in Air and Water*, H. Bleckmann, J. Mogdans, and S. L. Coombs, Eds., 8th ed. Berlin, Germany: Springer, 2014, pp. 521–562.

[16] M. Kruusmaa *et al.*, "FILOSE for svenning: A flow sensing bioinspired robot," *IEEE Robot. Automat. Mag.*, vol. 21, no. 3, pp. 51–62, Sep. 2014.

[17] Q. Hu, C. Wei, Y. Liu, and Z. Zhao, "A review of biomimetic artificial lateral line detection technology for unmanned underwater vehicles," in *Intelligent Robotics and Applications*, vol. 11741, H. Yu, J. Liu, L. Liu, Z. Ju, Y. Liu, and D. Zhou, Eds. Cham, Switzerland: Springer, 2019, pp. 501–516.

[18] G. Liu, M. Wang, A. Wang, S. Wang, T. Yang, R. Malekian, and Z. Li, "Research on flow field perception based on artificial lateral line sensor system," *Sensors*, vol. 18, no. 3, p. 838, Mar. 2018.

[19] A. T. Klein, F. Kaldenbach, A. Rüter, and H. Bleckmann, "What we can learn from artificial lateral line sensor arrays," in *The Effects of Noise on Aquatic Life II*, vol. 875, A. N. Popper and A. Hawkins, Eds. New York, NY, USA: Springer, 2016, pp. 539–545.

[20] A. B. Sichert, R. Bamler, and J. L. van Hemmen, "Hydrodynamic object recognition: When multipoles count," *Phys. Rev. Lett.*, vol. 102, no. 5, 2009, Art. no. 058104.

[21] B. Colvert, M. Alsalman, and E. Kanso, "Classifying vortex wakes using neural networks," *Bioinspiration Biomimetics*, vol. 13, no. 2, Jan. 2018, Art. no. 025003.

[22] N. Strokina, J.-K. Kamarainen, J. A. Tuhtan, J. F. Fuentes-Perez, and M. Kruusmaa, "Joint estimation of bulk flow velocity and angle using a lateral line probe," *IEEE Trans. Instrum. Meas.*, vol. 65, no. 3, pp. 601–613, Mar. 2016.

[23] V. I. Fernandez, A. Maertens, F. M. Yaul, J. Dahl, J. H. Lang, and M. S. Triantafyllou, "Lateral-line-inspired sensor arrays for navigation and object identification," *Mar. Technol. Soc. J.*, vol. 45, no. 4, pp. 130–146, Jul. 2011.

[24] B. J. Wolf, J. A. S. Morton, W. N. Macpherson, and S. M. Van Netten, "Bio-inspired all-optical artificial neuromast for 2D flow sensing," *Bioinspiration Biomimetics*, vol. 13, no. 2, Jan. 2018, Art. no. 026013.

[25] I. Guyon and A. Elisseeff, "An introduction to variable and feature selection," *J. Mach. Learn. Res.*, vol. 3, pp. 1157–1182, Jan. 2003.

[26] G.-B. Huang, Q.-Y. Zhu, and C.-K. Siew, "Extreme learning machine: Theory and applications," *Neurocomputing*, vol. 70, nos. 1–3, pp. 489–501, Dec. 2006.

[27] S. Coombs and S. van Netten, "The hydrodynamics and structural mechanics of the lateral line system," *Fish Physiol.*, vol. 23, pp. 103–139, 2005.

[28] A. J. Kalmijn, "Hydrodynamic and acoustic field detection," in *Sensory Biology of Aquatic Animals*. New York, NY, USA: Springer, 1988, pp. 83–130.

[29] S. Coombs and R. A. Conley, "Dipole source localization by mottled sculpin. I. Approach strategies," *J. Comparative Physiol. A: Sensory, Neural, Behav. Physiol.*, vol. 180, no. 4, pp. 387–399, Mar. 1997.

[30] B. J. Wolf and S. M. Van Netten, "Training submerged source detection for a 2D fluid flow sensor array with extreme learning machines," in *Proc. 11th Int. Conf. Mach. Vis. (ICMV)*, Mar. 2019, Art. no. 1104126.

[31] G. G. Stokes, *On the Effect of the Internal Friction of Fluids on the Motion of Pendulums*, vol. 9. Cambridge, U.K.: Pitt Press, 1851.

[32] H. Lamb, *Hydrodynamics*. Cambridge, U.K.: Cambridge Univ. Press, 1932.
[33] A. T. Abdulsadda and X. Tan, "Nonlinear estimation-based dipole source localization for artificial lateral line systems," *Bioinspiration Biomimetics*, vol. 8, no. 2, Mar. 2013, Art. no. 026005.
[34] R. Bouffanais, G. D. Weymouth, and D. K. P. Yue, "Hydrodynamic object recognition using pressure sensing," *Proc. Roy. Soc. A, Math., Phys. Eng. Sci.*, vol. 467, no. 2125, pp. 19–38, Jan. 2011.
[35] S. Pandya, Y. Yang, D. L. Jones, J. Engel, and C. Liu, "Multisensor processing algorithms for underwater dipole localization and tracking using MEMS artificial lateral-line sensors," *EURASIP J. Adv. Signal Process.*, vol. 2006, pp. 1–9, 2006.
[36] M. Alsalman, B. Colvert, and E. Kanso, "Training bioinspired sensors to classify flows," *Bioinspiration Biomimetics*, vol. 14, no. 1, Nov. 2018, Art. no. 016009.
[37] S. Lakkam, B. T. Balamurali, and R. Bouffanais, "Hydrodynamic object identification with artificial neural models," *Sci. Rep.*, vol. 9, no. 1, 2019, Art. no. 11242.
[38] A. Ristolainen, K. Kalev, J. A. Tuhtan, A. Kuusik, and M. Kruusmaa, "Hydromorphological classification using synchronous pressure and inertial sensing," *IEEE Trans. Geosci. Remote Sens.*, vol. 56, no. 6, pp. 3222–3232, Jun. 2018.
[39] D. Fan, G. Jodin, T. R. Consi, L. Bonfiglio, Y. Ma, L. R. Keyes, G. E. Karniadakis, and M. S. Triantafyllou, "A robotic intelligent towing tank for learning complex fluid-structure dynamics," *Sci. Robot.*, vol. 4, no. 36, Nov. 2019, Art. no. eaay5063.
[40] L. H. Boulogne, B. J. Wolf, M. A. Wiering, and S. M. Van Netten, "Performance of neural networks for localizing moving objects with an artificial lateral line," *Bioinspiration Biomimetics*, vol. 12, no. 5, Jul. 2017, Art. no. 056009.
[41] *Specification for Octave-Band and Fractional-Octave-Band Analog and Digital Filters*, Amer. Nat. Standards Inst., New York, NY, USA, 2004.
[42] D. Xu, Z. Lv, H. Zeng, H. Bessaih, and B. Sun, "Sensor placement optimization in the artificial lateral line using optimal weight analysis combining feature distance and variance evaluation," *ISA Trans.*, vol. 86, pp. 110–121, Mar. 2019.
[43] M. Robnik-Šikonja and I. Kononenko, "Theoretical and empirical analysis of ReliefF and RReliefF," *Mach. Learn.*, vol. 53, nos. 1–2, pp. 23–69, Oct. 2003.
[44] L. Yu and H. Liu, "Feature selection for high-dimensional data: A fast correlation-based filter solution," in *Proc. 20th Int. Conf. Mach. Learn. (ICML)*, 2003, pp. 856–863.
[45] H. Peng, F. Long, and C. Ding, "Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 8, pp. 1226–1238, Aug. 2005.
[46] C. Ding and H. Peng, "Minimum redundancy feature selection from microarray gene expression data," *J. Bioinform. Comput. Biol.*, vol. 3, no. 2, pp. 185–205, Apr. 2005.
[47] J. Tang, S. Alelyani, and H. Liu, "Feature selection for classification: A review," in *Data Classification: Algorithms and Applications*. Boca Raton, FL, USA: CRC Press, 2014, p. 37.
[48] B. Wolf and S. van Netten, "2D-sensitive flow sensor array data set for five different objects," Univ. Groningen, Groningen, The Netherlands, Tech. Rep., 2019. [Online]. Available: https://hdl.handle.net/10411/BWQMTQ, doi: 10411/BWQMTQ.

BEN J. WOLF received the M.Sc. degree in artificial intelligence from the University of Groningen, The Netherlands, in 2016, where he is currently pursuing the Ph.D. degree in hydrodynamic imaging with artificial intelligence, within the context of the EC funded Lakhsmi consortium.

His research interests include machine learning, neural networks, robotics, and hydrodynamic sensing.

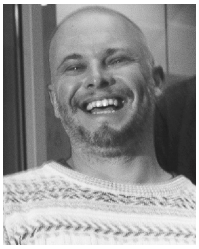

PRIMOZ PIRIH received the Ph.D. degree for his research on insect vision and coloration from the University of Groningen, The Netherlands, in 2011.

He is currently an Assistant Professor of physiology with the University of Ljubljana, Slovenia, where he is also developing light sources based on light emitting diodes for use in microscopy and in multi-spectral stimulation of animal and human visual systems. His research interests include lateral line sensing, and the evolution of color and polarization vision in insects.

MAARJA KRUUSMAA is currently a Professor of biorobotics and the Head of the Center for Biorobotics, Tallinn University of Technology, Tallinn, Estonia. She coordinated the EC funded Lakhsmi Project (sensors for large scale hydrodynamic imaging) under which the results presented in this article where achieved. Her research interests include flow sensing, underwater robots, bioinspired robot locomotion, and experimental fluid dynamics.

SIETSE M. VAN NETTEN is currently a Senior Lecturer with the Bernoulli Institute of Mathematics, Computing Science and Artificial Intelligence, University of Groningen, The Netherlands. He has been investigating the biophysics of the mechanosensory lateral line system over the last three decades. Results inspired the initiation of the all-optical artificial lateral line system, as described in the present article, which was further developed in the EC funded Lakhsmi Project.