

Multifidelity digital twin for real-time monitoring of structural dynamics in aquaculture net cages

Eirini Katsidoniotaki^{1,*}, Biao Su², Eleni Kelasidi², and Themistoklis P. Sapsis^{1,*}

¹Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02138, USA

²Department of Aquaculture Technology, SINTEF Ocean, Trondheim, Norway

*eirka289@mit.edu

*sapsis@mit.edu

ABSTRACT

As the global population grows and climate change intensifies, sustainable food production is critical. Marine aquaculture offers a viable solution, providing a sustainable protein source. However, the industry's expansion requires novel technologies for remote management and autonomous operations. Digital twin technology can advance the aquaculture industry, but its adoption has been limited. Fish net cages, which are flexible floating structures, are critical yet vulnerable components of aquaculture farms. Exposed to harsh and dynamic marine environments, the cages experience significant loads and risk damage, leading to fish escapes, environmental impacts, and financial losses. We propose a multifidelity surrogate modeling framework for integration into a digital twin for real-time monitoring of aquaculture net cage structural dynamics under stochastic marine conditions. Central to this framework is the nonlinear autoregressive Gaussian process method, which learns complex, nonlinear cross-correlations between models of varying fidelity. It combines low-fidelity simulation data with a small set of high-fidelity field sensor measurements, which offer the real dynamics but are costly and spatially sparse. Validated at the SINTEF ACE fish farm in Norway, our digital twin receives online metocean data and accurately predicts net cage displacements and mooring line loads, aligning closely with field measurements. The proposed framework is beneficial where application-specific data are scarce, offering rapid predictions and real-time system representation. The developed digital twin prevents potential damages by assessing structural integrity and facilitates remote operations with unmanned underwater vehicles. Our work also compares Gaussian processes and graph convolutional networks for predicting net cage deformation, highlighting the latter's effectiveness in complex structural applications.

Introduction

As the global population grows and the effects of climate change escalate, concerns about future food security and human nutrition increase, creating an urgent demand to enhance food production¹⁻³. Fish and fishery products have become vital in diets worldwide, particularly in developing countries⁴, but unsustainable fishing practices threaten natural fish populations⁵. Marine aquaculture emerges as a vital solution³, offering a sustainable protein source that meets the rising global food demands while serving as a response to climate change. Fish farms have lower carbon emissions⁶ compared to other animal proteins and support a diverse, resilient food system. Sustainable growth of aquaculture requires further reducing the carbon footprint of fish farms through strategic farm placement, efficient feed use, minimizing food waste and fish mortality rates, and the adoption of novel technologies to enhance autonomy in challenging fish farm operations. As the industry matures and operations move offshore to harsh and uncertain environments⁷, advancing real-time monitoring, operational planning, and decision-making practices is essential for ensuring efficiency, safety, cost-effectiveness, and sustainability^{8,9}.

Digital twins have great success in various sectors, yet its adoption in aquaculture has been limited, partly because it is a comparatively young industry⁹. Digital twin technology can significantly advance aquaculture by enabling fish farm remote monitoring, autonomous operations, management, and predictive analytics^{8,9}. Integrating real-time observations from in-situ monitoring equipment alongside realistic models within a digital twin framework can provide more detailed and accurate insights into the dynamics of fish farms. Implementing a digital twin typically involves several components^{8,9}, such as assessing fish health and behavior¹⁰, tracking environmental conditions¹¹, and monitoring the structural dynamics of fish net cages¹². Our study primarily focuses on the latter component.

Fish net cages are flexible floating structures anchored to the seabed, serving as critical yet vulnerable components of aquaculture farms. Exposed to the harsh and dynamic marine environment, these cages experience significant loads and are at risk of damage that can lead to fish escapes, resulting in negative environmental impacts and substantial financial losses¹³. Influenced by waves and currents, the net cages are prone to large motions and deformations¹⁴ due to their elastic structure.

These deformations reduce the available space for fish, causing decreased oxygen levels and increased stress¹⁵, which negatively impacts fish growth and survival, leading to higher mortality rates and significant food waste. The development of a digital twin capable to real-time monitor the net cage dynamics can serve as an early warning system for fish farm operators, helping to prevent significant reductions in cage volume. Additionally, a digital twin is essential for assessing the structural integrity and preventing potential catastrophic damages, while it can be used to facilitate the autonomous inspection and repair operations with unmanned underwater vehicles (UUVs) in fish farms. For example, when autonomously controlling a remotely operated vehicle (ROV) within the fish cage to assess potential damage, prior knowledge of the cage's shape is crucial for effective path planning and collision avoidance with the net¹⁶.

Unlike more rigid marine structures such as floating wind turbine foundations or vessels, the elastic and complex geometry of aquaculture net cages—composed of millions of slender twines—poses unique challenges in predicting their dynamics¹⁷. Modeling these net cages is a multifaceted task that involves solving sub-problems: one focuses on the interaction between the fluid and a rigid moored-floating structure, and the other addresses the flexible porous net. Accurately depicting the deformation of these cages requires detailed information on the displacements of each twine, necessitating extensive data and computational effort. While high-fidelity (HF) numerical models like CFD simulations provide accurate estimations of structural dynamics¹⁷, they demand vast computational resources, making them unsuitable for real-time forecasting or rapid hindcasting tasks in digital twins. Lower fidelity (LF) models^{18–21} are more accessible and less computationally demanding but often fail to capture the real system accurately, particularly under non-linear effects from extreme waves and fluid-structure interactions, thus limiting their integration into digital twins. Conversely, field sensor measurements offer the most accurate depiction of net cage dynamics but are constrained by high sensor costs and the risk of data transmission loss¹². Additionally, sensors typically provide localized information, resulting in significant spatial gaps or high sparsity. These challenges underscore the need for innovative approaches to effectively monitor net cage system dynamics in real-time.

Current research is focused on innovative models for precise real-time monitoring of net cage dynamics that can be integrated into digital twins. Su et al.¹⁴ introduced a real-time numerical simulation model that uses in-situ, real-time positioning sensor data to dynamically adjust model inputs—specifically, the magnitude and direction of the current—based on the difference between simulated and actual net positions. Building on this, Su et al. further developed their model into a physics-based digital twin¹² that incorporates in-situ sensor data to accurately simulate real-time cage responses, net deformations, and mooring loads. Although successfully tested at a full-scale aquaculture site, this model faces challenges such as high computational costs and potential inaccuracies due to inconsistent or faulty sensor data caused by harsh environmental conditions, hardware limitations, and unexpected interference, making it less suited for real-time applications. These issues highlight the need for digital twins using data-driven and machine learning (ML)-based surrogate models^{22,23}.

ML-based surrogate models for predicting the dynamic behavior of fish net cages remain an under-explored research area. Some studies utilized artificial neural networks (ANNs) for net damage detection^{24–26}, yet research on predicting the structural dynamics of aquaculture net cages remains sparse. Zhao et al.²⁷ developed a backpropagation ANN with one hidden layer, using training data from LF numerical simulations, to correlate ocean waves with net cage structural responses, including maximum tension in mooring lines, minimum effective volume ratio of the cage, and maximum stress on the floating collar. Although their model outputs the volume ratio, it does not capture the precise topology of the deformed net under wave excitation. To the best of our knowledge, no surrogate models have been developed to predict the entire net cage deformation, which can be beneficial for preventing catastrophic events and supporting autonomous fish farm operations. This underscores the need for more advanced ML-based surrogate models, built on high-fidelity data, that can accurately predict the complex dynamics of flexible net cages. To support effective real-time monitoring, these models should be capable of integrating into on-the-fly digital twins, which means their computational demands need to be limited.

Creating surrogate models that accurately represent complex systems often encounters the challenge of acquiring HF, application-specific data. Multifidelity data assimilation²⁸ involves methods that merge information about the same underlying truth obtained through multiple models or observation operators at different fidelity levels, resulting in a model that best represents the system's state. Among these methods, optimal interpolation is a widely used data assimilation technique; however, it may exhibit large discrepancies with the truth²⁹. To this end, multifidelity surrogate modeling³⁰ aims to leverage machine learning methods with computationally inexpensive low-fidelity (LF) data and scarce high-fidelity (HF) data to accurately predict the quantities of interest. Despite their limitations, LF data can reveal system trends and patterns, providing valuable information to complement the limited HF data used for model training. Multifidelity methods^{30–36} for constructing surrogate models aim to achieve effective data fusion from various fidelity levels, enabling strong generalization performance of data-driven models in regions where HF data are limited. This approach is particularly beneficial for various digital twin applications^{37–40}, enabling accurate, real-time system representation.

We propose a framework specifically designed for seamless integration into digital twins, enabling real-time monitoring and visualization of aquaculture net cage structural dynamics under stochastic marine conditions. The framework is built on a multifidelity surrogate modeling method that integrates data from various fidelity sources representing the fish farm

environment. Specifically, the framework utilizes the Nonlinear Autoregressive Gaussian Process (NARGP) method³², which is capable of learning complex, nonlinear, and space-dependent cross-correlations between models of varying fidelity. Due to the reduced algorithmic complexity and computational cost of the NARGP, this framework offers rapid predictions essential for real-time representation of the current system's state. The resulting multifidelity digital twin enables effective monitoring of structural integrity and the detection of irregularities, functioning as an alarm system to guide necessary actions and prevent damage that could lead to fish escapes or mortalities. Moreover, it informs decision-making and enhances operational planning by supporting the control of autonomous underwater operations in fish farms, such as net damage inspections and net cleaning tasks with ROVs. This framework significantly enhances efficiency, automation, safety, and cost savings by reducing the need for additional sensors and personnel visits, which are particularly challenging in harsh marine environments. Figure 1 illustrates the proposed multifidelity digital twin.

Although NARGP has been validated in benchmark problems, its application in real-world scenarios has not been widely explored. To demonstrate the effectiveness of our framework, we validated it on a full-scale net cage at the SINTEF ACE fish farm in Norway, a leading laboratory facility for developing and testing new aquaculture technologies. Data for this study were sourced from the FhSim software¹⁹ and field sensor measurements. Leveraging both HF and LF data, our framework ensures that predictions closely align with real observations. This framework accurately predicts the net cage's displacement at specific locations and the mooring line loads under current excitations. Additionally, in our study, we explore different ML methods for developing surrogate models to predict the entire net cage deformation. We compare the Gaussian process (GP) with a prior principal component analysis (PCA) step and graph convolutional networks (GCNs)^{41,42}. While GCNs are widely used in life, physical, and material sciences, our research demonstrates their effectiveness in handling complex structural dynamics⁴³.

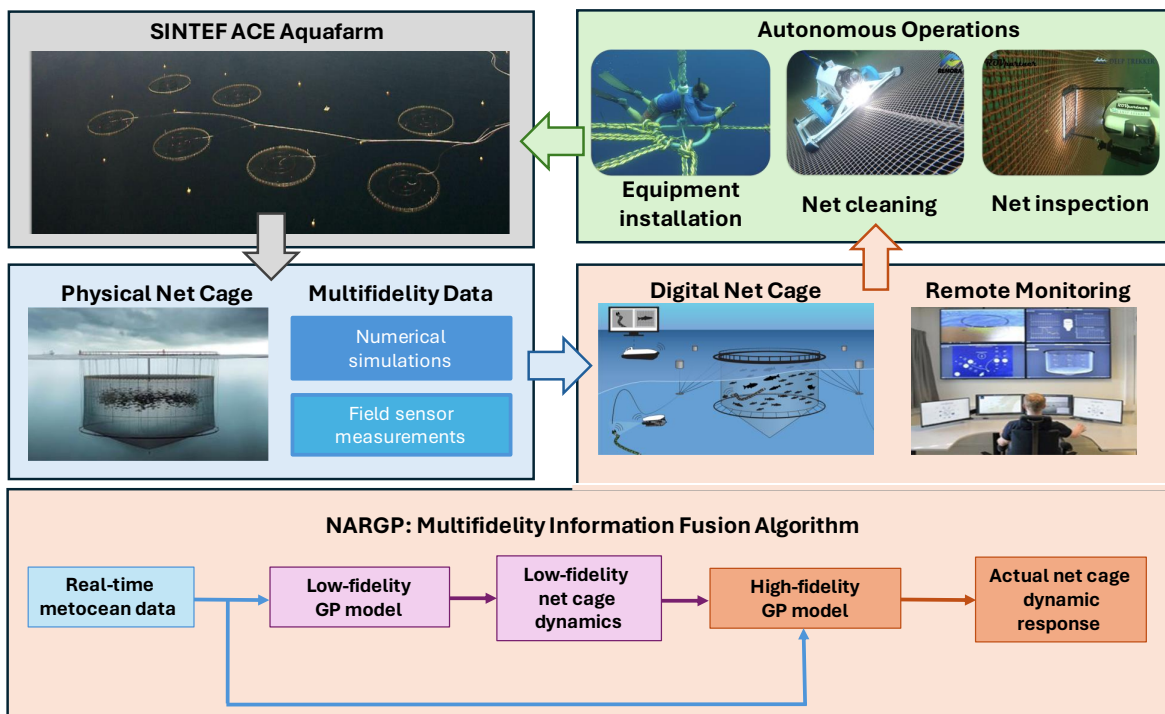


Figure 1. Real-time monitoring and remote fish farm management using digital twin technology. The structural response of a net cage to dynamic marine environments is modeled using data from varying fidelity sources, including low-fidelity numerical simulations and high-fidelity field sensor measurements. A digital twin of the physical net cage is developed using the proposed multifidelity framework, central to which is the Nonlinear Autoregressive Gaussian Process (NARGP) method³². This method synergistically combines low-fidelity model realizations with a limited set of high-fidelity observations. The nonlinear autoregressive scheme effectively learns complex nonlinear cross-correlations between datasets of differing fidelity, which is essential for practical applications. The multifidelity digital twin accurately predicts the net cage's structural response under dynamic marine conditions, closely mirroring real-world behavior. The digital twin facilitates remote fish farm management, serving as an early warning system to prevent catastrophic damage, informing decision-making, and supporting autonomous operations such as ROV path planning for net inspection and repair, net cleaning, and equipment installation.

Results

Our goal is to develop a framework for a digital twin that enables real-time monitoring and visualization of aquaculture net cage dynamics under stochastic sea conditions, ensuring efficiency, safety, and cost-effectiveness while minimizing environmental impacts. Machine-learned surrogate modeling is fundamental in order to create the map between the marine conditions and the net cage dynamic responses. Accurate prediction of real-world behaviors relies on comprehensive, high-quality data, but acquiring such data is often challenging or cost-prohibitive. Building reliable models becomes complex when only a limited number of HF observations, such as direct sensor measurements, are available. This limitation often leads to the use of models that represent the real system based on more accessible data, such as those from simplified numerical models. However, a significant gap typically exists between these simulated outcomes and actual sensor readings. Additionally, it is crucial that this model can be trained and executed rapidly to function as an on-the-fly digital twin. Therefore, employing less complex methods is essential to achieve fast performance.

We propose a framework that utilizes data from various fidelity sources to learn the system dynamics, producing outputs that closely match real-world observations. At the core of this framework is the NARGP method, as described by Perdikaris et al. (2017)³², which refines low-fidelity numerical simulation solutions to more closely match high-fidelity sensor data by learning the nonlinear correlations between these two data sources using data-efficient multifidelity information fusion algorithms and the simple GP methods. The NARGP method employs a sequential training process across each fidelity level. Detailed description of the NARGP is provided in the section Methods. For our case study, which is detailed in the next section, the model is designed with two fidelity levels but is scalable to include additional layers from different data fidelities if necessary. The trained NARGP model is implemented in the digital twin following the workflow illustrated in Figure 1. It receives real-time metocean data, and the GP model, trained on the low-fidelity numerical data, predicts the net cage dynamics. These low-fidelity predictions, along with the metocean data, are then input into the recursive GP surrogate model. This model has been trained using the GP posterior mean derived from the low-fidelity model and a few high-fidelity data, and it is able to provide predictions that accurately correspond to the actual net cage response.

The proposed multifidelity framework enables the use of simple, low-cost models that may have lower accuracy, and significantly enhances their precision by incorporating a small set of high-fidelity observations. By leveraging the nonlinear cross-correlations between low- and high-fidelity data through machine learning, this approach achieves substantial computational efficiency. This method allows us to tackle complex problems that would be infeasible to address using only high-fidelity data, which are often difficult to obtain due to cost and other limitations.

Case study for SINTEF ACE fish farm

The multifidelity digital twin for real-time monitoring of net cage structural dynamics is built and validated for a fish farm at [SINTEF ACE](#), which is a full-scale laboratory facility designed to develop and test new aquaculture technologies. All relevant data used in this study correspond to this farm.

Data Collection

Low-fidelity data. FhSim is a software platform and framework for mathematical modeling and numerical simulation, with a focus on marine applications, providing a time-domain representation of complex system dynamics. In this study, FhSim is used to model the dynamic response of a full-scale net cage to marine environment. Figure 2 (a) displays the numerical net cage in the FhSim environment, detailing its critical components. The numerical net cage has the same dimensions as the real cage, featuring a 50 m diameter, 18 m depth, and a 0.21 solidity ratio, attached by 12 bridle lines to the mooring frame. Although the actual net cage comprises thousands of twines, for numerical simplification, the net cage is discretized into 321 nodes, as shown in Figure 2 (b). In FhSim simulations, the displacements of each of the 321 nodes under wave and current excitation are calculated, fully capturing the net cage's deformation, as shown in Figure 2 (c).

For the purpose of this study, we assessed the net cage's dynamic behavior under 1,000 sea conditions covering a broad spectrum of wave and current conditions. The environmental parameters included current velocities ranging from 0 to 1 m/s, significant wave heights from 0 to 3 m, peak wave periods from 0 to 8.66 s, and wave and current directions from 0 to 360 degrees. Each simulation modeled the net cage interaction with the sea conditions over a 30-minute period, providing as output the x , y , and z displacements for each of the 321 nodes and the loads on the 12 mooring lines. We focused on estimating the average value of the quantities of interest under each examined sea condition.

FhSim was previously used by Su et al. (2023)¹² to simulate net cage dynamics under measured wave and current conditions in the fish farm. The direct solution of the FhSim model was found to underestimate net deformations, and this was partially attributed to the increased net solidity due to biofouling, as well as the possible alterations of the current from the measuring point (i.e., 400 m away). The authors¹² developed and integrated an adaptive model that makes connections to the environmental

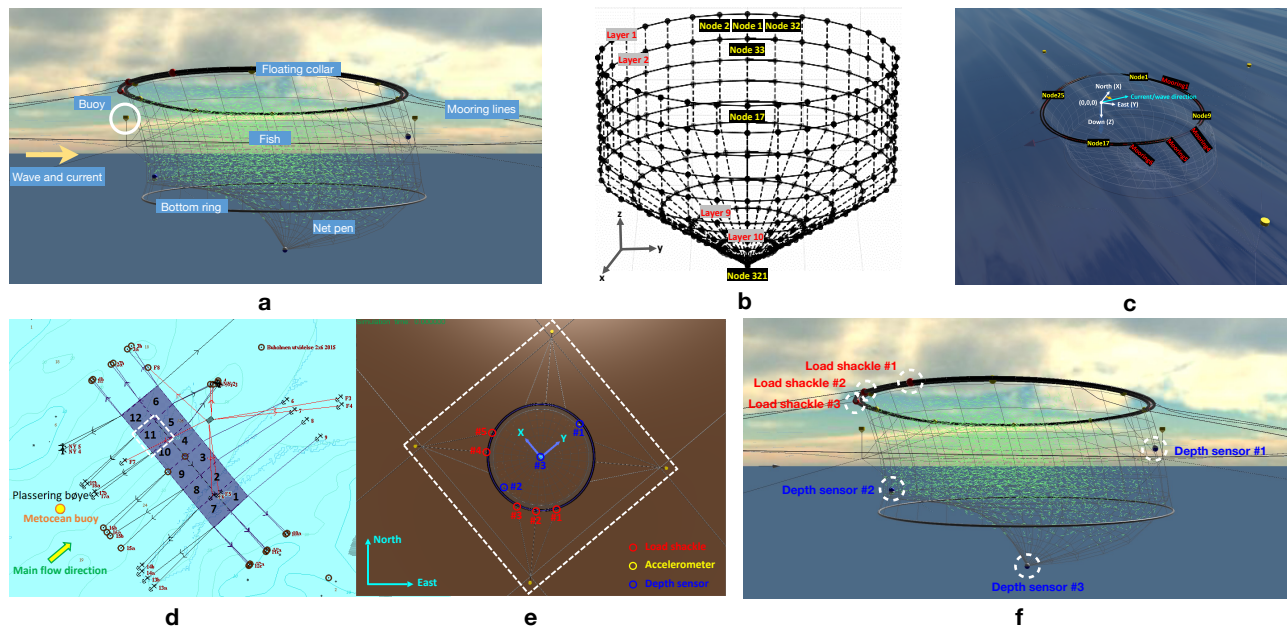


Figure 2. Net cage numerical representation and field sensor placement. (a) Detailed view of net cage’s critical components. (b) The net cage is discretized into 321 nodes for FhSim numerical simulations. It consists of 10 layers, with the first 32 nodes forming the first layer. Node 321 is at the bottom, connecting all nodes of the 10th layer. (c) The net cage in the FhSim simulation environment, showing deformation under wave and current excitation. (d) Top view of the fish farm highlighting net cage 11, which accommodates the sensors. A buoy located 400 meters from the net cage measures waves and currents, with the main flow direction indicated. (e) Close-up of net cage 11 showing sensor locations: 5 load shackles measuring mooring line forces and 3 depth sensors measuring net displacement. (f) Side view of the net cage illustrating sensor locations along its depth.

condition inputs so that the calculated structural dynamics match observations. The FhSim’s direct solution aims to predict the steady state of net cage’s dynamics, rather than all the transient responses; therefore, it is viewed in our study as low-fidelity source of data. Apart from the solution accuracy, conducting FhSim simulations requires computational resources and time. Specifically, in our study, a 30-minute simulation time took approximately 10 minutes of CPU time (utilizing 56 cores, 2.2 GHz, and 128 GB RAM), with an additional 20 minutes required for processing and storing time-series data for all mooring lines and nodes. This makes it impractical for direct implementation in a digital twin for real-time monitoring, where immediate knowledge of the net cage response is essential.

High-fidelity data. Field sensor measurements capture the actual response of the system and are considered high-fidelity data. As shown in Figure 2 (d), a metocean buoy, positioned 400 meters away from the fish farm, was equipped with sensors to monitor the sea conditions, specifically measuring incoming waves and currents. Sensors were placed in one of the net cages at SINTEF fish farm, specifically in cage number 11 (Figure 2 (d)). The sensors included five load shackles (#1–5 in Figure 2(e)) installed between the collar and the mooring lines on the upstream side, where the prevailing water current originated, to measure mooring line loads. Depth sensors were installed at three different layers: 7 meters deep, at the boundary of the cylindrical part (15 meters), and at the bottom (31 meters) of the net pen (#1–3 in Figures 2 (e) and (f)). These depth sensors measured the vertical displacement of the net cage at their respective positions, providing critical insights into the net cage deformation. The data were collected using a wireless sensor network and communication system over a monitoring period of slightly less than two months, from January 18 to March 5, 2020. This period captures conditions representative of a relatively harsh time of year. Data points are available at a high temporal resolution of 2-second intervals, providing comprehensive insights into the variations and dynamics of the measured quantities. Continuous measurements were recorded by the metocean buoy and the three depth sensors covering the entire period from January 18 to March 5, 2020. However, load shackle data were only available for a total of 36 hours—12 hours on January 23 and 24 hours on February 22, highlighting the risk of data loss that field sensors may experience. Detailed information about the field deployment and sensor measurements is available in the publication by Su et al. (2023)¹².

Digital twin workflow

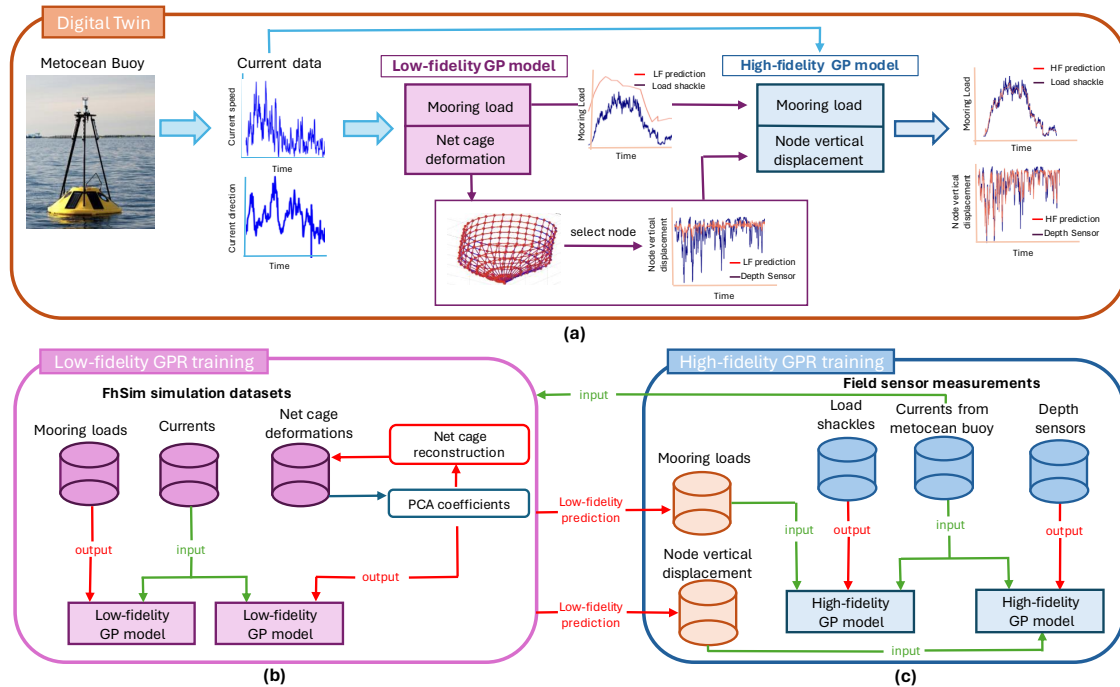


Figure 3. Multifidelity digital twin for SINTEF ACE fish farm. (a) The digital twin’s workflow is illustrated. The digital twin receives real-time measurements of current speed and direction from a metocean buoy equipped with sensors. These measurements serve as inputs to the low-fidelity Gaussian Process (GP) models, which predict the mooring line loads and the net cage deformation. The low-fidelity predictions, represented by the GP posterior mean values, along with the metocean buoy’s current measurements, are subsequently fed into high-fidelity GP models which refine these predictions, achieving greater accuracy that closely matches the actual observations from load shackles and depth sensors. (b) The training stage of the low-fidelity GP models using data from FhSim simulations is described. A standard GP model is trained to map current characteristics to mooring line loads. For net cage deformation, dimension reduction using PCA is performed first due to the extensive data required to describe the flexible net topology. A standard GP model maps current characteristics to PCA coefficients, followed by a reconstruction stage that outputs the predicted net deformation. These low-fidelity GP models are capable of substituting FhSim simulations. (c) The training stage of the high-fidelity GP models based on multifidelity data is illustrated. Limited field sensor measurements randomly selected from January 18 to March 5, 2020, are used for training, including current measurements from the metocean buoy, mooring loads from load shackles, and net cage vertical displacements at depth sensor positions. The current measurements are input into the low-fidelity GP models, and the resulting low-fidelity predictions, combined with the current measurements, serve as inputs for the high-fidelity models. These models are trained to provide predictions that closely match load shackle and depth sensor measurements.

Figure 3 (a) illustrates the digital twin workflow developed for the SINTEF ACE fish farm. A metocean buoy equipped with sensors provides real-time measurements of current speed and direction which serve as inputs to the NARGP multifidelity surrogate model. Initially, the low-fidelity GP models use the current data to predict the loads on the mooring lines and the overall deformation of the net cage under these conditions. To estimate the net cage displacement at specific locations, the user defines these locations (or nodes) on the net cage and obtains their vertical displacement. The low-fidelity predictions of mooring line force and node vertical displacement, along with the measured current speed and direction, are then fed into high-fidelity GP models. These models refine the low-fidelity predictions, providing outputs that closely match the real structural dynamics measured by load shackles and depth sensors. This workflow is completed in a few seconds due to the algorithmic simplicity and low computational cost of the NARGP method. The proposed digital twin enables real-time monitoring of net cage structural dynamics, supports autonomous operations, and informs decision-making without relying on real-time sensors from the net cage, as is currently the case¹². This approach reduces the need for extensive sensor deployment, which is often expensive and unreliable due to issues like data loss.

Figure 3 (b) shows the training stage of the low-fidelity GP surrogate models using data from FhSim simulations, described

in the "Data Collection/Low-fidelity" section. Using the standard GP method⁴¹, we build surrogate models to map current speed and direction to the mooring line load. To create the surrogate model that maps ocean currents to net cage deformation, dimension reduction using PCA is required as a preliminary step, as detailed in the "Methods/Net cage displacement data dimension reduction using PCA" section. This is necessary due to the extensive data describing net cage deformation (x , y , z displacements for each of the 321 nodes). From the PCA dimension reduction step, we found that three PCA coefficients were sufficient and this allowed us to develop a standard GP model to map current characteristics to these PCA coefficients. Once the predicted coefficients are obtained, a reconstruction stage follows to determine the net cage deformation under the incoming measured current, as described in the "Methods/Machine learning functional relationships between current and PCA coefficients" section. Having the entire geometry of the deformed net cage allows us to extract the vertical displacement at the locations where the depth sensors are placed. The low-fidelity GP models developed at this stage can fully substitute FhSim simulations, with their accuracy well-documented by Katsidoniotaki et al. (2024)⁴⁴. This ensures that if the load shackles and depth sensors change position or more sensors are added in the future, retraining the low-fidelity GP models will not be necessary.

Figure 3 (c) outlines the steps for training the high-fidelity GP models following the NARGP method, as described in the "Methods/NARGP" section. The high-fidelity training data consist of field sensor measurements, including current measurements from the metocean buoy, mooring loads from load shackles, and net cage vertical displacements from depth sensors, as detailed in the "Data Collection/High-Fidelity" section. The current measurements are input into the low-fidelity GP models, and the predicted GP posterior mean, combined with the current measurements, serve as inputs for the high-fidelity models. The high-fidelity GP models are trained on a limited number of datasets obtained from the period January 18 to March 5, 2020, to predict quantities of interest that match the load shackle and depth sensor measurements. Specifically, while the depth sensor measurements cover the entire period, only a small number of datasets corresponding to a few minutes are used to train the model. For the load shackle measurements, available only for 36 hours, datasets corresponding to a few seconds are used for training. This approach demonstrates how, with limited high-fidelity datasets, we can develop a multifidelity model that provides predictions representative of real observations.

Predicting net cage dynamics

Mooring line loads. Figure 4 showcases the proposed multifidelity framework's proficiency in predicting mooring line loads under current conditions, as measured by the metocean buoy, and compares these predictions with field measurements obtained from load shackles. The figure presents predictions for load shackles #1 (top row), #2 (middle row), and #5 (bottom row), with similar patterns observed for other mooring lines. Validation was conducted using previously unseen data from the entire 36-hour load shackle measurement period, excluding the training datasets. The 2D histograms in Figure 4 are scatter plots with a color scale depicting the distribution and density of prediction accuracy from the low-fidelity and multifidelity GP models. The color scale indicates the log count density of data points, with lighter colors representing areas with a higher concentration of data points. Proximity to the dashed diagonal line ($y = x$) signifies prediction accuracy, while the red line represents the best-fit trend. The scatter plots for the low-fidelity GP model (Figure 4 (a), (d), (g)) reveal a broader spread and a significant deviation of the red best-fit line from the dashed diagonal line, indicating less accurate predictions. Conversely, the multifidelity GP model plots (Figure 4 (b), (e), (h)) exhibit a higher concentration of lighter colors near the dashed line, with the red best-fit line aligning more closely with the diagonal, indicating improved prediction accuracy. The color gradient helps identify trends and outliers; consistent, lighter colors along the dashed line suggest the model captures the overall trend well, whereas significant deviations indicate less accurate predictions. The Mean Absolute Error (MAE) quantifies prediction accuracy, calculated as the average of the absolute differences between the predicted and actual load measurements. The multifidelity model demonstrates a significant improvement in MAE values compared to the low-fidelity model. Figure 4 (c), (f), (i) provides a temporal view, displaying predictions over time based on the validation data from the 36-hour measurement period. The low-fidelity GP posterior mean deviates notably from the actual measurements, whereas the multifidelity GP posterior aligns closely with the real observations. The temporal plots reveal an oscillatory behavior in the load shackle measurements not fully captured by the multifidelity model, explaining some deviations from the dashed line in the middle column's scatter plot. Nonetheless, the multifidelity model accurately follows the overall trend of the observations.

Vertical displacement at specific points within the net cage. Figure 5 illustrates the predictive performance of the multifidelity framework in estimating net cage displacement at three distinct locations, as measured by depth sensors #1 (top row), #2 (middle row), and #3 (bottom row). The validation dataset spans from January 18 to March 5, 2020, excluding the training periods. Figure 5 (a), (d), (g) presents scatter plots contrasting the predictions of the low-fidelity GP model with actual sensor measurements. The broad dispersion of data points and the notable deviation of the red best-fit line from the dashed diagonal line underscore the model's limited accuracy. In contrast, Figure 5 (b), (e), (h) displays scatter plots for the multifidelity GP model. Here, data points cluster more closely around the dashed line, and the red best-fit line aligns more precisely with the diagonal, indicating enhanced predictive accuracy. Despite the inherent complexities and dynamic motion of the flexible

net cage, the multifidelity model's predictions, though still dispersed, show improved accuracy, as evidenced by the color scale: darker points represent outliers, while lighter points closer to the dashed line signify higher accuracy. The MAE values annotated on these plots quantitatively highlight the multifidelity model's superiority over the low-fidelity model. Figure 5 (c), (f), (i) provides a temporal comparison of predicted and actual net cage displacements over a specific period (January 18 to January 27, 2020), juxtaposing the low-fidelity GP and multifidelity GP posterior means against the actual sensor data. The low-fidelity GP model (blue line) exhibits substantial deviations from the actual measurements (green line), particularly for depth sensors #1 and #3. Conversely, the multifidelity GP model (red line) significantly mitigates these discrepancies, closely mirroring the actual measurements and capturing the underlying trend with greater fidelity.

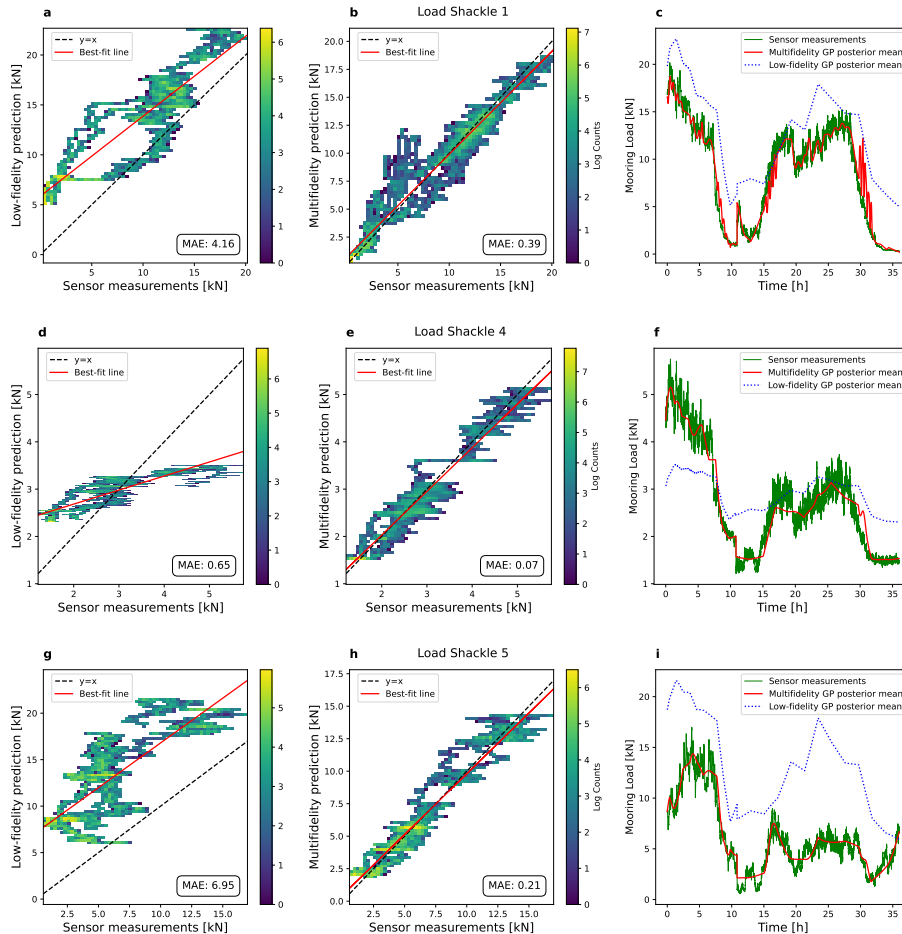


Figure 4. Multifidelity framework predictions for mooring line loads. Scatter plots (a), (d), (g) show the low-fidelity GP posterior mean against load shackle measurements. Color scale indicates log count density of data points, with lighter colors representing higher concentrations. The dashed diagonal line signifies perfect prediction accuracy; the red line is the best-fit trend. Low-fidelity GP predictions have broad spread and the red line presents significant deviation from the dashed line, indicating less accurate predictions. Scatter plots (b), (e), (h) show the multifidelity GP model predictions against sensor measurements. Multifidelity GP model has higher concentration of lighter colors near the dashed line, with the red line closely aligning with it, indicating improved accuracy. Plots (c), (f), (i) display predictions over a 36-hour period, excluding the data used for training. Low-fidelity GP posterior mean deviates significantly from actual measurements. Multifidelity GP posterior aligns closely with real observations, capturing the overall trend despite some oscillatory behavior deviations. Mean Absolute Error (MAE) values indicate a significant improvement in prediction accuracy with the multifidelity model compared to the low-fidelity model.

Flexible net cage deformation via surrogate modeling

Fish farm net cages are flexible structures that tend to follow rather than resist water motions, therefore they present significant deformations. These deformations can impact fish survival and the structural integrity of the cages, with any failure potentially

leading to fish escapes. This underscores the importance of monitoring net cage deformations to maintain a comprehensive understanding of fish farm conditions. Due to their inherent flexibility, monitoring the entire deformation of net cages using available sensors is challenging. Depth sensors can be mounted strategically on the net structure thereby obtaining a 3D position of these points in the net, which in turn are used to extrapolate the full net structure using mathematical models⁹. A similar approach was implemented by Su et al.¹² for the development of a digital twin for real-time monitoring of aquaculture net cage systems, where depth sensor data were assimilated into FhSim simulations to represent the actual net cage system. Input properties for the simulations were adapted to ensure the numerical simulation outputs fit the sensor-measured values. However, this approach depends on sensors transmitting position information, with sensors being expensive and prone to data transmission loss. Additionally, it requires substantial computational time to process the data and calculate the correct net cage deformation. These reasons make this approach less suitable for real-time monitoring and immediate assessment of the system's state.

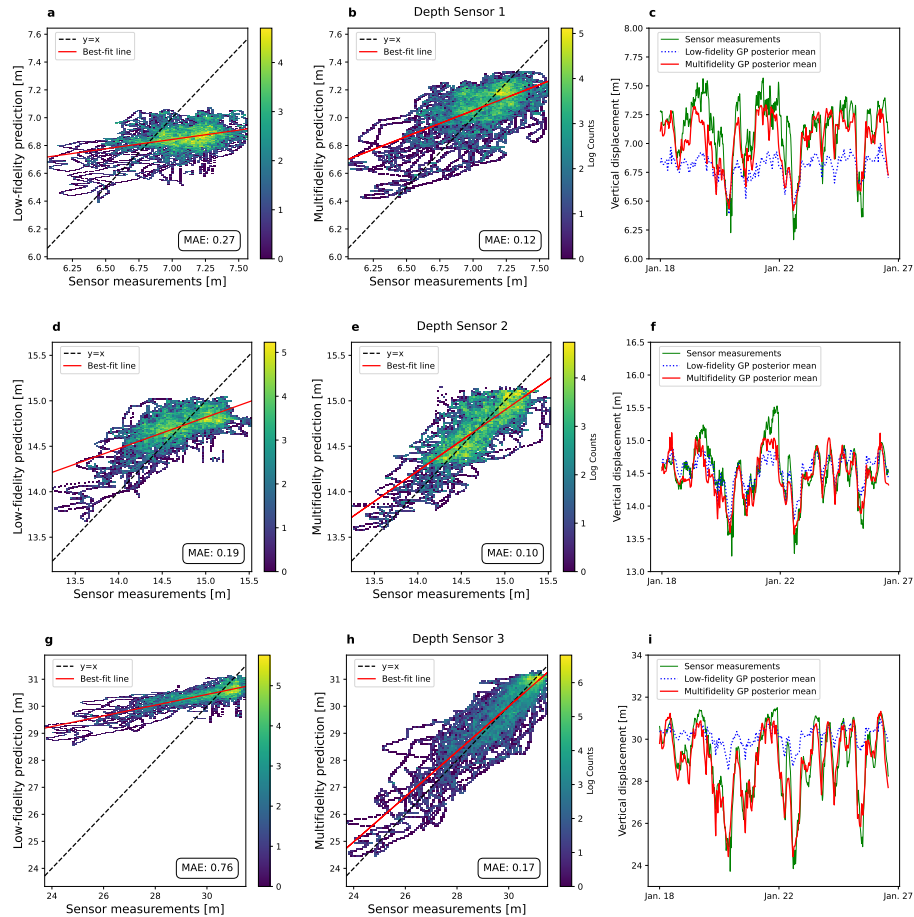


Figure 5. Multifidelity framework predictions for net cage displacement. Scatter plots (a), (d), (g) contrast the predictions from the low-fidelity GP model with the depth sensor measurements. The broad dispersion of data points and the notable deviation of the red best-fit line from the dashed diagonal line underscore the model's limited accuracy. Color scale indicates log count density of data points, with lighter colors representing higher concentrations. Scatter plots (b), (e), (h) for the multifidelity GP model. Data points cluster more closely around the dashed line, and the red best-fit line aligns more precisely with the diagonal, indicating enhanced predictive accuracy. Despite the inherent complexities and dynamic motion of the flexible net cage, the multifidelity model's predictions show improved accuracy. Darker points represent outliers, while lighter points closer to the dashed line signify higher accuracy. The MAE values annotated on these plots quantitatively highlight the multifidelity model's superiority over the low-fidelity model. Plots (c), (f), (i) show the temporal comparison of predicted and actual net cage displacements over a specific period (January 18 to January 27, 2020). The low-fidelity GP model (blue line) exhibits substantial deviations from the actual measurements (green line), particularly for depth sensors #1 and #3. The multifidelity GP model (red line) significantly mitigates these discrepancies, closely mirroring the actual measurements and capturing the underlying trend with greater fidelity.

Our study previously introduced the multifidelity digital twin, which accurately reflects the real net cage response to the dynamic marine environment. As a result, this digital twin eliminates the reliance on real-time data measured by sensors placed in the cage. Additionally, as illustrated in Figure 3 (b), we developed surrogate models, using the standard GP method with a prior PCA, that map the net cage deformation with the currents. The accuracy of these models, previously presented by Katsidoniotaki et al.⁴⁴, suggests that they can replace FhSim simulations and significantly reduce computational costs in digital twin applications. However, surrogate models inherently introduce some degree of prediction error, depending on the method used to build them. This error can impact autonomous operations or decision-making where good accuracy is required. In this study, we compare the GP-PCA surrogate model with a surrogate model built using GCNs to predict net cage deformation topology under varying sea conditions. We chose the GCN since the net cage resembles graph-structured data. The cage is discretized into 321 nodes, which are connected in the manner shown in Figure 2 (b). The GCN receives node features and the topology of their connections and can predict the displacement of each node under dynamic marine conditions. In the section "Methods/Graph Convolutional Networks for Net Cage Deformation," further details about the GCN model are provided. Figure 6 compares the predicted net cage deformation from the GCN model (red) with the deformation estimated by FhSim simulations (blue) under both mild (test index 801) and harsher sea conditions (test index 31).

To evaluate the prediction error of each method, we consider the example of autonomous navigation of Unmanned Underwater Vehicles (UUVs), which require real-time information about the surrounding environment where they navigate. During autonomous operations, a UUV such as a Remotely Operated Vehicle (ROV) should maintain a distance of at least 1 meter from the net cage to ensure collision-free motion. Therefore, it is imperative to keep the prediction error regarding net cage deformation within acceptable limits for safe operations. The MAE metric was employed to quantify the accuracy of net cage deformation predictions across three spatial directions over a range of test indices. As depicted in Figure 6, the majority of data points for both methods lie below the red line at MAE = 1 meter, indicating a satisfactory level of accuracy for most test indices. Notably, higher MAE values are observed in test cases involving higher current velocities, where the net cage exhibits larger deformations, highlighting the challenges in modeling more extreme conditions. As summarized in the final subplot in Figure 6, the GP-PCA model demonstrates higher MAE values, suggesting lower prediction accuracy compared to the GCN model. The wider spread of the GP-PCA error bars in the x and z dimensions notably suggests greater variability in its predictions. This greater error in the GP-PCA model may be attributed to dual sources of error: the PCA dimension reduction, which entails some loss of information, and the inherent errors in the GP model itself.

Discussion

In this study, we present a framework that leverages multifidelity surrogate modeling capabilities, designed to be integrated into a digital twin for real-time monitoring of structural dynamics. Unlike traditional digital twins that rely on physics-based simulations—requiring significant computational resources and thus being suboptimal for real-time monitoring—our approach offers several key advantages. Our proposed framework requires less computational cost during both the training and utilization stages, allowing it to operate on the fly and provide rapid predictions. Furthermore, our framework effectively addresses a common challenge in real-world applications: the limited availability and access to real-world data for developing realistic predictive surrogate models. By seamlessly integrating low-fidelity numerical data with high-fidelity field sensor measurements, our model bridges the gap between simulation outcomes and real-world observations. This integration enhances the accuracy and reliability of predictions, making the digital twin highly suitable for real-time monitoring applications, even in scenarios where real-world data is scarce.

The NARGP method, central to our framework, belongs to a class of multifidelity information fusion algorithms. An important feature of the NARGP method is the nonlinear autoregressive scheme that enables robust correlation among different fidelity data sources, addressing the challenges posed by nonlinear relationships. While NARGP has been validated against benchmark problems, it has not yet been validated for a real application. The NARGP method performs well in our application, providing good prediction accuracy by significantly correcting the low-fidelity solution to match the real solution quickly and at a low computational cost.

The framework successfully validated for the test case of SINTEF ACE fish farm, using low-fidelity numerical simulation data and high-fidelity field sensor measurements. For the net cage's vertical displacement at specific locations, the low-fidelity GP prediction deviates significantly from the depth sensor measurements. This discrepancy is likely due to the pronounced nonlinear and complex phenomena associated with the hydrodynamic effects due to the flexible net cage's interaction with the wave and currents, which cannot be fully captured by numerical simulations. However, the NARGP method effectively corrects these deviations, resulting in a final prediction that closely aligns with the depth sensors. Our framework was developed and validated using high-fidelity data from field sensor measurements obtained over a period slightly less than two months,

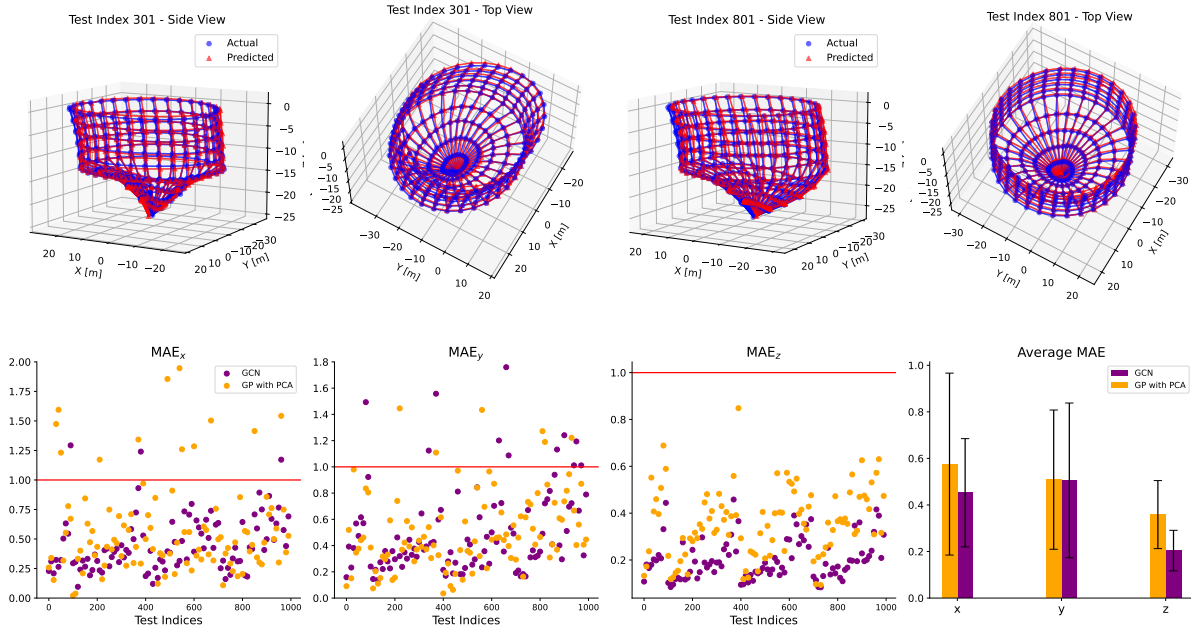


Figure 6. Machine-learned surrogate models for predicting net cage deformation. (a) Comparison of net cage deformation predictions under harsh (301) and mild (801) current excitations by Graph Convolutional Networks (GCNs) with FhSim simulation solutions. (b) Evaluation of net cage deformation prediction accuracy using MAE metric, comparing predictions of each node displacement in the x , y , z directions with the FhSim solution. Two surrogate modeling methods are compared: (i) GP surrogate with prior PCA, and (ii) GCNs. The red line at $MAE = 1$ serves as a reference point, indicating the threshold beyond which prediction error significantly impacts operations such as UUV control. The surrogates are evaluated across 100 test cases. The final plot summarizes the MAE, showing the average magnitude of errors in predictions for each axis, with error bars representing the variability in the error.

specifically from January 18 to March 5, 2020. Although this period covered harsh sea conditions, ideally, measurements should span the entire year to account for seasonal variations. In the future, if more field sensor data become available, the high-fidelity surrogate model can be retrained using new samples that cover chosen from the total available measurements, further improving the framework’s accuracy and robustness.

Accurate predictions of mooring loads and net cage deformation are crucial for evaluating structural integrity, informing decisions on predictive maintenance, and preventing structural damage that could lead to fish mortality or escape. This digital twin also enables autonomous operations in fish farms, such as inspecting nets and mooring lines for damages, irregularities, biofouling conditions, and net cleaning⁴⁵. These operations are currently performed using UUVs such as ROVs. For the autonomous navigation of these systems, it is essential to plan and control the motion of the ROV’s motion to ensure collision-free movement in highly dynamic environments^{45,46}. Since the ROV operates inside the net cage, real-time monitoring of net cage deformation is essential to avoid collisions.

Additionally, we offer a solution for the challenging task of predictive modeling of net cage deformation due to the structure’s flexibility. Knowing the exact net cage deformation is critical for fish welfare, net cage structural integrity, and overall fish farm operations. Our study demonstrates the capability of surrogate modeling to predict the deformation of the entire net cage under diverse marine conditions. We showcase the effectiveness of GCNs in assimilating the graph-structured net cage data, achieving better prediction accuracy compared to the GP method with PCA dimension reduction. This study is one of the few to utilize GCNs for structural deformation problems, highlighting the potential for GCNs to be expanded beyond their current applications in fields such as drug discovery, biology, transportation, and computer vision.

In conclusion, in this study we introduce a solution aided by machine learning and multifidelity data assimilation for developing digital twins for the remote fish farm management, aiming to advance aquaculture industry with new practices. While our focus has been on predicting net cage deformation and mooring line loads, the multifidelity framework is generalizable for learning any stochastic nonlinear behavior of other quantities of interest in structural problems and other aquaculture matters, such as

predicting fish behavior and environmental conditions in fish farms. The proposed digital twin solution improves efficiency and safety, prevents damages, reduces costs, and mitigates adverse environmental impacts. These factors are essential for the industry's expansion and the push towards offshore waters to meet rising food demands and secure food supply amid climate change.

Methods

Nonlinear autoregressive GP regression (NARGP)

The nonlinear autoregressive GP regression denoted by NARGP³² is a class of multifidelity nonlinear information fusion algorithms that enables accurate inference of quantities of interest by synergistically combining realizations of low-fidelity models with a small set of high-fidelity observations, capable of learning complex nonlinear and space-dependent cross-correlations between models of variable fidelity. Our research adopts the NARGP method in response to the observed nonlinear dynamics between our low and high-fidelity datasets.

We have s levels of information sources (fidelities) producing outputs $y_t(\mathbf{x}_t)$, at locations $\mathbf{x}_t \in D_t \subseteq \mathbb{R}^d$, with $t = 1, \dots, s$. In our application we have two fidelity levels, $s = 2$, and two input parameters, $d = 2$, i.e. current velocity and direction. We can organize the available datasets by increasing fidelity as $D_t = \{\mathbf{x}_t, \mathbf{y}_t\}$, while the datasets have nested structure, i.e. $D_2 \subseteq D_1$. This assumption implies that the training inputs of the higher fidelity level need to be a subset of the training inputs of the lower fidelity level. In NARGP³² method, the nonlinear autoregressive scheme reads as:

$$f_t(\mathbf{x}) = g_t(\mathbf{x}, f_{*_{t-1}}(\mathbf{x})) \quad (1)$$

where f_t is the GP modeling the data at fidelity level t , the $f_{*_{t-1}}$ is the GP posterior from the previous inference level $t - 1$, and the function g_t is modeled as a Gaussian prior, $g_t \sim \mathcal{G}(\mathbf{f}_t | \mathbf{0}, \mathbf{k}_{t_g})$, with a covariance kernel that composes as

$$k_{t_g} = k_{t_\rho}(\mathbf{x}, \mathbf{x}'; \theta_{t_\rho}) \cdot k_{t_f}(f_{*_{t-1}}(\mathbf{x}), f_{*_{t-1}}(\mathbf{x}'); \theta_{t_f}) + k_{t_\delta}(\mathbf{x}, \mathbf{x}'; \theta_{t_\delta}) \quad (2)$$

The structure of the kernel k_{t_g} reveals the effect of the deep representation encoded in Eq. 1. The $g_t(\mathbf{x}, f_{*_{t-1}}(\mathbf{x}))$ projects the lower fidelity posterior $f_{*_{t-1}}$ onto a $(d + 1)$ -dimensional latent manifold, that jointly relates the input space and the outputs of the lower fidelity level to the output of the higher fidelity model, from which we can infer a smooth mapping that recovers the high-fidelity response f_t . This allows to capture general nonlinear, non-functional and space-dependent cross-correlations between low- and high-fidelity data. In Eq. 2, k_{t_ρ} , k_{t_f} , and k_{t_δ} are valid covariance functions and θ_{t_ρ} , θ_{t_f} , θ_{t_δ} denote their hyperparameters, which can be easily learnt from the data $\{\mathbf{x}_t, \mathbf{y}_t\}$ via the maximum likelihood estimation procedure followed by the standard GP⁴¹ using the kernel k_{t_g} . This approach requires the estimation of $(2d + 3)$ hyperparameters assuming that all kernels account for directional anisotropy in each dimension using automatic relevance determination (ARD) weights. The kernel functions are chosen to have the squared exponential form with ARD weights, i.e.

$$k_t(x, x'; \theta_t) = \sigma_t^2 \exp \left(-\frac{1}{2} \sum_{i=1}^d w_{i,t} (x_i - x'_i)^2 \right) \quad (3)$$

where σ_t^2 is a variance parameters and $(w_{i,t})_{i=1}^d$ are the ARD weights corresponding to fidelity level t . These weights allow for a continues 'blend' of the contributions of each individual dimension in \mathbf{x}_t as well as the posterior predictions of the previous fidelity level $f_{*_{t-1}}$, and they are learnt directly from the data when inferring f_t .

The first level of the proposed recursive scheme corresponds to a standard GP regression problem trained on the lowest fidelity data $\{\mathbf{x}_1, \mathbf{y}_1\}$, and therefore, the predictive posterior distribution is defined by standard's Gaussian mean and covariance using the kernel function $k_1(\mathbf{x}_1, \mathbf{x}'_1; \theta_1)$. Figure 3 provides a summary of the NARGP workflow as implemented in this study.

Net cage displacement data dimension reduction using PCA

PCA is a statistical technique used in data analysis for dimensionality reduction. It identifies the directions (or principal components) that maximize the variance in a dataset, thereby preserving essential information with fewer variables. In our study, we apply PCA to the dataset obtained from FhSim numerical simulations, which provide the x_i, y_i, z_i displacements of the nodes into which the net cage is discretized, where $i = 1, \dots, 321$. We have this data for each of the $s = 1000$ environmental conditions examined. Specifically, for each environmental condition, FhSim simulates the net cage interaction with waves and currents over a 30-minute period, outputting the $x_i(t), y_i(t), z_i(t)$ node displacements. The 30-min simulation time for all the $s = 1000$ environmental conditions corresponds to a total of t_n time steps. All the available data from these simulations are

integrated into the matrix \mathbf{M} .

$$\mathbf{M} = \begin{bmatrix} x_1(t_1) & \dots & x_1(t_n) \\ y_1(t_1) & \dots & y_1(t_n) \\ z_1(t_1) & \dots & z_1(t_n) \\ \dots & \dots & \dots \\ x_{321}(t_1) & \dots & x_{321}(t_n) \\ y_{321}(t_1) & \dots & y_{321}(t_n) \\ z_{321}(t_1) & \dots & z_{321}(t_n) \end{bmatrix}$$

Using this data matrix, we perform eigenanalysis to calculate the eigenvectors, ϕ , and eigenvalues, λ , respectively.

$$\mathbf{M}\mathbf{M}^T \phi = \lambda \phi \quad (4)$$

The eigenvalues and their corresponding eigenvectors are sorted in descending order of the eigenvalues. The eigenvector with the highest eigenvalue is the first principal component of the dataset and accounts for the highest variance, with each subsequent component capturing progressively less variance than its predecessor. We determine the number of principal components to retain based on the cumulative explained variance ratio, which translates into how many components k we should keep out of the c total components to better describe the data. This sum yields the total proportion of the dataset's variance that is accounted for by the first k principal components.

$$\text{Cumulative Variance Explained} = \sum_{p=1}^k \frac{\lambda_p}{\sum_{j=1}^c \lambda_j} \quad (5)$$

A common approach is to choose the smallest possible number of principal components, k , that explain a substantial portion of the variance. However, the choice of k is based on a predetermined threshold which reflects the desired level of variance explanation: Cumulative Variance Explained \geq Threshold. In our case, setting the threshold at 93% results in the retention of $k = 3$ principal components.

Next, we project the original data, \mathbf{M} , onto the selected principal components to transform them into the new feature subspace, by multiplying the original matrix by the first k selected eigenvectors, to get the \mathbf{B} matrix of the transformed data which can be referred as PCA coefficients.

$$\mathbf{B} = \mathbf{M} \cdot \phi \quad (6)$$

The resulting \mathbf{B} matrix consists of 3 rows, each corresponding to one of the principal components, and t_n columns each one corresponds to a time step state of the net cage deformation. Since we are interested in estimating the average shape of the net cage deformation for each of the $s = 1000$ scenarios, we estimate the mean principle component for each scenario, resulting in a matrix $\bar{\mathbf{B}}$, which can be referred as mean PCA coefficients. This matrix has $k = 3$ rows and $s = 1000$ columns.

$$\bar{\mathbf{B}} = \begin{bmatrix} \bar{b}_{1,1} & \dots & \bar{b}_{1,s} \\ \bar{b}_{2,1} & \dots & \bar{b}_{2,s} \\ \bar{b}_{3,1} & \dots & \bar{b}_{3,s} \end{bmatrix}$$

To reconstruct the average net cage deformation for each of the total $s = 1000$ environmental conditions, we multiple the eigenvector ϕ with the principal components.

$$\bar{\mathbf{M}} = \sum_{h=1}^s \sum_{j=1}^k \phi_j \cdot \bar{b}_{jh} \quad (7)$$

The matrix that shows the average net cage displacement for each of the examined environmental conditions is:

$$\bar{\mathbf{M}} = \begin{bmatrix} \bar{x}_{1,1} & \dots & \bar{x}_{1,s} \\ \bar{y}_{1,1} & \dots & \bar{y}_{1,s} \\ \bar{z}_{1,1} & \dots & \bar{z}_{1,s} \\ \dots & \dots & \dots \\ \bar{x}_{321,1} & \dots & \bar{x}_{321,s} \\ \bar{y}_{321,1} & \dots & \bar{y}_{321,s} \\ \bar{z}_{321,1} & \dots & \bar{z}_{321,s} \end{bmatrix}$$

Machine learning functional relationships between current and PCA coefficients

To develop the low-fidelity component of the NARG methodology, which predicts net cage deformation that corresponds to FhSim solution, we create a standard GP model to learn the functional relationship between current speed and the net cage deformation. Specifically, the GP model maps the current characteristics to the PCA coefficients, $\bar{b}_1, \bar{b}_2, \bar{b}_3$. Then, we utilize the inner sum of Equation 7 ($\sum_{j=1}^k \phi_j \cdot \bar{b}_j$) to reconstruct the net cage deformation for an individual environmental condition. In other words, we obtain the predicted x_i, y_i, z_i , where $i = 1, \dots, 321$. Similar approaches—incorporating PCA for dimensionality reduction and the development of surrogate models using machine learning methods—have been widely applied across various disciplines due to their efficacy in feature extraction and data simplification^{37,47–50}. In our case, performing the dimension reduction step helps reduce the complexity of the data describing the net cage deformation, making it easier to use machine learning methods to create the mapping.

Graph convolutional networks for net cage deformation

Graph neural networks (GNNs) are mathematical models that can learn functions over graphs and are a learning approach for building predictive models on graph-structured data. Graphs differ from regular data in that they have a structure that neural networks must respect. Among the various types of GNNs, the graph convolutional networks (GCNs) have emerged as the most broadly applied model. GCNs are innovative due to their ability generalize the operation of convolutional neural networks to graphs, enabling the network to learn from the graph’s topology and node features to make predictions.

Representing net cage graph. The net cage of our application can be viewed as graph-structured data, discretized into 321 nodes. We build a GCN model for the net cage based on the work of Kipf and Welling⁴². The net cage can be viewed as a graph G consisting of $N = 321$ nodes connected through the edges, where v_i and v_j represent nodes i and j . The edge features are represented by the adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ which encodes connections between the N nodes, represents the graph connectivity which is a square matrix where each element A_{ij} specifies the presence or absence of an edge from node i to node j in the graph. In other words, a non-zero element A_{ij} implies a connection from node i to node j , and a zero indicates no direct connection. In our net cage application, the adjacent matrix \mathbf{A} captures the node connection shown in Figure 2 (b); the net cage consists of 10 layers, each with 32 nodes arranged in a circular pattern. Nodes within each layer are interconnected to form a circle, and each node is vertically connected to the node directly below it in the subsequent layer. This creates a consistent connection pattern between layers. Additionally, the final node (node 321) is connected to all nodes in the 10th layer, providing a central connection point. The $\mathbf{X} \in \mathbb{R}^{L \times F}$ is the node feature matrix with F being the number of the feature assigned to each node. In our case, the \mathbf{X} contains the features for each node which in our case are the x_o, y_o, z_o coordinates of each node at rest and the sea conditions (i.e. current speed and direction). The graph is supposed to predict the matrix $\mathbf{Y} \in \mathbb{R}^{L \times O}$, where O is the number of output features per node. The matrix \mathbf{Y} gives the x, y, z coordinates of each node under the interaction of the cage with the sea condition.

Message aggregation and update. The GCN layers are based on the formulation by Kipf and Welling (2017)⁴². At each layer, the graph convolution takes both the adjacency matrix \mathbf{A} and the node features from the previous layer $\mathbf{H}^l \in \mathbb{R}^{L \times F_l}$ and outputs the node features for the next layer $\mathbf{H}^{l+1} \in \mathbb{R}^{L \times F_{l+1}}$, with F_l and F_{l+1} are the node features dimensions for layers l and $l + 1$, while $\mathbf{H}^0 = \mathbf{X}$ the input node feature matrix, following the formulation:

$$\mathbf{H}^{l+1} = \sigma(\tilde{\mathbf{D}}^{-0.5} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-0.5} \mathbf{H}^l \mathbf{W}^l) \quad (8)$$

where $\sigma(\cdot)$ is the activation function, $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}_L$ is the adjacency matrix with added self-connections represented by the identity matrix $\mathbf{I}_L \in \mathbb{R}^{L \times L}$, $\tilde{\mathbf{D}}$ is the diagonal degree matrix of $\tilde{\mathbf{A}}$, and $\mathbf{W}^l \in \mathbb{R}^{F_l \times F_{l+1}}$ is a trainable weight matrix for layer $l + 1$. In our application, the data passes through three graph convolution layers. Each layer updates the node features by aggregating information from its neighbor nodes according to the graph structure. The final output represents the predicted node features which are the predicted displacements (x, y, z) for each node. Having several layers enables to further refine the node features by aggregating information from the graph structure and learning more complex patterns. The deformation of net-cage nodes is a highly non-linear process influenced, the ability of Swish activation function to introduce non-linearities is particularly beneficial, which is defined as:

$$\sigma(x) = x \cdot \frac{1}{1 + e^{-x}} \quad (9)$$

Model training. A notable point in our model is the implementation of a custom loss function that assigns different weights to the errors of each node by emphasizing regions of the net-cage that are more prone to displacement. At the node weights initialization stage nodes within layers 6 to 10, which are observed to undergo significant deformation, are assigned increased weights, with the highest weight allocated to the 321st node. The loss function computes the mean squared error (MSE) for

each nodes prediction:

$$\text{Loss} = \frac{1}{N} \sum_{i=1}^N w_i \cdot \left(\frac{1}{F} \sum_{j=1}^F (\hat{y}_{ij} - y_{ij})^2 \right) \quad (10)$$

given $\hat{\mathbf{Y}} = [\hat{\mathbf{y}}_1, \hat{\mathbf{y}}_2, \dots, \hat{\mathbf{y}}_N]$: Predicted displacements (matrix of shape $N \times F$), $\mathbf{Y} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N]$: True displacements (matrix of shape $N \times F$), $\mathbf{w} = [w_1, w_2, \dots, w_N]$: Node weights (vector of length N). For the training process we use backpropagation to compute gradients and gradient descent optimiser to update the parameters for each layer.

Data availability statement

Code will become available to a dedicated depository after the paper is accepted. Correspondence and requests for code materials should be addressed to Eirini Katsidoniotaki or Themistoklis P. Sapsis. The data that support the findings of this study are available from SINTEF Ocean AS but restrictions apply to the availability of these data, which were used under license for the current study, and so are not publicly available. Data are however available from the authors upon reasonable request and with permission of SINTEF Ocean AS.

References

1. Johan Bergenas. The Other Global Food Crisis. <https://foreignpolicy.com/2023/09/19/global-food-crisis-fishing-blue-foods-conflict-water-ocean-climate-resources/>.
2. United Nations. The World's Food Supply is Made Insecure by Climate Change. <https://www.un.org/en/academic-impact/worlds-food-supply-made-insecure-climate-change/>.
3. World Economic Forum. Aquaculture: why the world needs a new wave of food production. <https://www.weforum.org/agenda/2022/01/aquaculture-agriculture-food-systems/> (2022).
4. Han, K., Leem, K., Choi, Y. R. & Chung, K. What drives a country's fish consumption? market growth phase and the causal relations among fish consumption, production and income growth. *Fish. Res.* **254**, 106435, DOI: <https://doi.org/10.1016/j.fishres.2022.106435> (2022).
5. WWF. Overfishing. <https://www.worldwildlife.org/threats/overfishing/>.
6. Gephart, J. A. *et al.* Environmental performance of blue foods. *Nature* **597**, 360–365 (2021).
7. Bjelland, H. V. *et al.* Exposed aquaculture in norway. In *OCEANS 2015-MTS/IEEE Washington*, 1–10 (IEEE, 2015).
8. Ubina, N. A. *et al.* Digital twin-based intelligent fish farming with artificial intelligence internet of things (aiot). *Smart Agric. Technol.* **5**, 100285, DOI: <https://doi.org/10.1016/j.atech.2023.100285> (2023).
9. Føre, M. *et al.* Digital twins in intensive aquaculture—challenges, opportunities and future prospects. *Comput. Electron. Agric.* **218**, 108676, DOI: <https://doi.org/10.1016/j.compag.2024.108676> (2024).
10. Saad, A., Su, B. & Bjørnson, F. O. A web-based platform for efficient and robust simulation of aquaculture systems using integrated intelligent agents. *Procedia Comput. Sci.* **225**, 4560–4569, DOI: <https://doi.org/10.1016/j.procs.2023.10.454> (2023).
11. Huan, J., Li, H., Li, M. & Chen, B. Prediction of dissolved oxygen in aquaculture based on gradient boosting decision tree and long short-term memory network: A study of chang zhou fishery demonstration base, china. *Comput. Electron. Agric.* **175**, 105530, DOI: <https://doi.org/10.1016/j.compag.2020.105530> (2020).
12. Su, B. *et al.* Towards a holistic digital twin solution for real-time monitoring of aquaculture net cage systems. *Mar. Struct.* **91**, 103469, DOI: <https://doi.org/10.1016/j.marstruc.2023.103469> (2023).
13. Føre, H. M. & Thorvaldsen, T. Causal analysis of escape of atlantic salmon and rainbow trout from norwegian fish farms during 2010–2018. *Aquaculture* **532**, 736002, DOI: <https://doi.org/10.1016/j.aquaculture.2020.736002> (2021).
14. Su, B. *et al.* An integrated approach for monitoring structural deformation of aquaculture net cages. *Ocean. Eng.* **219**, 108424, DOI: <https://doi.org/10.1016/j.oceaneng.2020.108424> (2021).
15. Turnbull, J., Bell, A., Adams, C., Bron, J. & Huntingford, F. Stocking density and welfare of cage farmed atlantic salmon: application of a multivariate analysis. *Aquaculture* **243**, 121–132, DOI: <https://doi.org/10.1016/j.aquaculture.2004.09.022> (2005).
16. Amundsen, H. B., Xanthidis, M., Føre, M., Ohrem, S. J. & Kelasidi, E. Aquaculture field robotics: Applications, lessons learned and future prospects. *arXiv preprint arXiv:2404.12995* (2024).

17. Martin, T., Tsarau, A. & Bihs, H. A numerical framework for modelling the dynamics of open ocean aquaculture structures in viscous fluids. *Appl. Ocean. Res.* **106**, 102410, DOI: <https://doi.org/10.1016/j.apor.2020.102410> (2021).
18. Su, B. *et al.* A multipurpose framework for modelling and simulation of marine aquaculture systems. In *International conference on offshore mechanics and arctic engineering*, vol. 58837, V006T05A002 (American Society of Mechanical Engineers, 2019).
19. Reite, K.-J. *et al.* Fhsim—time domain simulation of marine systems. In *International Conference on Offshore Mechanics and Arctic Engineering*, vol. 45509, V08AT06A014 (American Society of Mechanical Engineers, 2014).
20. Kristiansen, T. & Faltinsen, O. M. Modelling of current loads on aquaculture net cages. *J. Fluids Struct.* **34**, 218–235, DOI: <https://doi.org/10.1016/j.jfluidstructs.2012.04.001> (2012).
21. Huang, C.-C., Tang, H.-J. & Liu, J.-Y. Dynamical analysis of net cage structures for marine aquaculture: Numerical simulation and model testing. *Aquac. Eng.* **35**, 258–270, DOI: <https://doi.org/10.1016/j.aquaeng.2006.03.003> (2006).
22. Liu, X., Han, X., Wang, H. & Liu, G. A modelling and updating approach of digital twin based on surrogate model to rapidly evaluate product performance. *The Int. J. Adv. Manuf. Technol.* **129**, 5059–5074, DOI: <https://doi.org/10.1007/s00170-023-12646-w> (2023).
23. Chakraborty, S., Adhikari, S. & Ganguli, R. The role of surrogate models in the development of digital twins of dynamic systems. *Appl. Math. Model.* **90**, 662–681, DOI: <https://doi.org/10.1016/j.apm.2020.09.037> (2021).
24. Bi, C.-W. *et al.* An efficient artificial neural network model to predict the structural failure of high-density polyethylene offshore net cages in typhoon waves. *Ocean. Eng.* **196**, 106793, DOI: <https://doi.org/10.1016/j.oceaneng.2019.106793> (2020).
25. Zhao, Y.-P., Lian, L., Bi, C.-W. & Xu, Z. Digital twin for rapid damage detection of a fixed net panel in the sea. *Comput. Electron. Agric.* **200**, 107247, DOI: <https://doi.org/10.1016/j.compag.2022.107247> (2022).
26. Zhang, Z., Gui, F., Qu, X. & Feng, D. Netting damage detection for marine aquaculture facilities based on improved mask r-cnn. *J. Mar. Sci. Eng.* **10**, 996, DOI: <https://doi.org/10.3390/jmse10070996> (2022).
27. Zhao, Y.-P., Bi, C.-W., Sun, X.-X. & Dong, G.-H. A prediction on structural stress and deformation of fish cage in waves using machine-learning method. *Aquac. Eng.* **85**, 15–21, DOI: <https://doi.org/10.1016/j.aquaeng.2019.01.003> (2019).
28. Popov, A. A. & Sandu, A. Multifidelity data assimilation for physical systems. *Data Assim. for Atmospheric, Ocean. Hydrol. Appl. (Vol. IV)* 43–67, DOI: https://doi.org/10.1007/978-3-030-77722-7_2 (2022).
29. Babae, H., Bastidas, C., Defilippo, M., Chryssostomidis, C. & Karniadakis, G. A multifidelity framework and uncertainty quantification for sea surface temperature in the massachusetts and cape cod bays. *Earth Space Sci.* **7**, e2019EA000954, DOI: <https://doi.org/10.1029/2019EA000954> (2020).
30. Conti, P. *et al.* Multi-fidelity reduced-order surrogate modelling. *Proc. Royal Soc. A* **480**, 20230655, DOI: <https://doi.org/10.1098/rspa.2023.0655> (2024).
31. Kennedy, M. C. & O’Hagan, A. Predicting the output from a complex computer code when fast approximations are available. *Biometrika* **87**, 1–13, DOI: <https://doi.org/10.1093/biomet/87.1.1> (2000).
32. Perdikaris, P., Raissi, M., Damianou, A., Lawrence, N. D. & Karniadakis, G. E. Nonlinear information fusion algorithms for data-efficient multi-fidelity modelling. *Proc. Royal Soc. A: Math. Phys. Eng. Sci.* **473**, 20160751, DOI: <https://doi.org/10.1098/rspa.2016.0751> (2017).
33. Cutajar, K., Pullin, M., Damianou, A., Lawrence, N. & González, J. Deep gaussian processes for multi-fidelity modeling. *arXiv preprint arXiv:1903.07320* (2019).
34. Meng, X. & Karniadakis, G. E. A composite neural network that learns from multi-fidelity data: Application to function approximation and inverse pde problems. *J. Comput. Phys.* **401**, 109020, DOI: <https://doi.org/10.1016/j.jcp.2019.109020> (2020).
35. Meng, X., Babae, H. & Karniadakis, G. E. Multi-fidelity bayesian neural networks: Algorithms and applications. *J. Comput. Phys.* **438**, 110361, DOI: <https://doi.org/10.1016/j.jcp.2021.110361> (2021).
36. Zhang, L., Wu, Y., Jiang, P., Choi, S.-K. & Zhou, Q. A multi-fidelity surrogate modeling approach for incorporating multiple non-hierarchical low-fidelity data. *Adv. Eng. Informatics* **51**, 101430, DOI: <https://doi.org/10.1016/j.aei.2021.101430> (2022).
37. Champenois, B. & Sapsis, T. Machine learning framework for the real-time reconstruction of regional 4d ocean temperature fields from historical reanalysis data and real-time satellite and buoy surface measurements. *Phys. D: Nonlinear Phenom.* **459**, 134026, DOI: <https://doi.org/10.1016/j.physd.2023.134026> (2024).

38. Lai, X. *et al.* Digital twin-based structural health monitoring by combining measurement and computational data: An aircraft wing example. *J. Manuf. Syst.* **69**, 76–90, DOI: <https://doi.org/10.1016/j.jmsy.2023.06.006> (2023).
39. Desai, A. S., Navaneeth, N., Adhikari, S. & Chakraborty, S. Enhanced multi-fidelity modeling for digital twin and uncertainty quantification. *Probabilistic Eng. Mech.* **74**, 103525, DOI: <https://doi.org/10.1016/j.probengmech.2023.103525> (2023).
40. Chetan, M., Yao, S. & Griffith, D. T. Multi-fidelity digital twin structural model for a sub-scale downwind wind turbine rotor blade. *Wind. Energy* **24**, 1368–1387, DOI: <https://doi.org/10.1002/we.2636> (2021).
41. Rasmussen, C. E., Williams, C. K. *et al.* *Gaussian processes for machine learning*, vol. 1 (Springer, 2006).
42. Kipf, T. N. & Welling, M. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
43. Li, X., Xu, L., Guo, H. & Yang, L. Application of graph convolutional neural networks combined with single-model decision-making fusion neural networks in structural damage recognition. *Sensors* **23**, 9327, DOI: <https://doi.org/10.3390/s23239327> (2023).
44. Katsidoniotaki, E., Su, B., Kelasidi, E. & Sapsis, T. Integrating machine learning for real-time structural monitoring of net cages. In *34th International Ocean and Polar Engineering Conference (ISOPE 2024)* (2024).
45. Kelasidi, E. & Svendsen, E. Robotics for sea-based fish farming. In *Encyclopedia of Smart Agriculture Technologies*, 1–20, DOI: https://doi.org/10.1007/978-3-031-24861-0_202 (Springer, 2023).
46. Kelasidi, E. *et al.* Autonomous monitoring and inspection operations with uavs in fish farms. *IFAC-PapersOnLine* **55**, 401–408, DOI: <https://doi.org/10.1016/j.ifacol.2022.10.461> (2022).
47. Rudy, S. H. & Sapsis, T. P. Prediction of intermittent fluctuations from surface pressure measurements on a turbulent airfoil. *AIAA J.* **60**, 4174–4190 (2022).
48. Guth, S. & Sapsis, T. P. Wave episode based gaussian process regression for extreme event statistics in ship dynamics: Between the scylla of karhunen–loève convergence and the charybdis of transient features. *Ocean. Eng.* **266**, 112633 (2022).
49. Katsidoniotaki, E., Guth, S., Götteman, M. & Sapsis, T. P. Reduced order modeling of wave energy systems via sequential bayesian experimental design and machine learning. (2023).
50. Guth, S., Katsidoniotaki, E. & Sapsis, T. P. Statistical modeling of fully nonlinear hydrodynamic loads on offshore wind turbine monopile foundations using wave episodes and targeted cfd simulations through active sampling. *Wind. Energy* **27**, 75–100 (2024).

Acknowledgements

This work was supported by the Knut and Alice Wallenberg Foundation Postdoctoral Scholarship Program at MIT - KAW 2022.0334, the Office of Naval Research (Grant no N00014-21-1-2357), by RACE internal funding (SINTEF Ocean AS) and by the Research Council of Norway (RCN) project: CHANGE (no. 313737). The experimental data utilized in this paper for validation of the framework were obtained from RACE internal funding and FLEXAQUA (no. 284361) project funded from RCN.

Author contributions statement

All authors conceived the idea. E.Kat. and T.P.S developed the methodology for the multifidelity digital twin framework. E.Kat. analysed the data, developed the code and implemented the framework. B.S. conducted the FhSim simulations. B.S. and E.Kel. provided the field sensor measurements. E.Kat. and T.P.S. analysed the results. B.S. and E.Kel. suggested applications for the digital twin’s practical implementation. E.Kat. wrote the original manuscript, with all authors contributing to editing and reviewing the manuscript.

Additional information

Competing interests statement

The authors declare no competing interests.

Legends

Figure 1: **Real-time monitoring and remote fish farm management using digital twin technology.** The structural response of a net cage to dynamic marine environments is modeled using data from varying fidelity sources, including low-fidelity numerical simulations and high-fidelity field sensor measurements. A digital twin of the physical net cage is developed using the proposed multifidelity framework, central to which is the Nonlinear Autoregressive Gaussian Process (NARGP) method³². This method synergistically combines low-fidelity model realizations with a limited set of high-fidelity observations. The nonlinear autoregressive scheme effectively learns complex nonlinear cross-correlations between datasets of differing fidelity, which is essential for practical applications. The multifidelity digital twin accurately predicts the net cage's structural response under dynamic marine conditions, closely mirroring real-world behavior. The digital twin facilitates remote fish farm management, serving as an early warning system to prevent catastrophic damage, informing decision-making, and supporting autonomous operations such as ROV path planning for net inspection and repair, net cleaning, and equipment installation.

Figure 2: **Net cage numerical representation and field sensor placement.** (a) Detailed view of net cage's critical components. (b) The net cage is discretized into 321 nodes for FhSim numerical simulations. It consists of 10 layers, with the first 32 nodes forming the first layer. Node 321 is at the bottom, connecting all nodes of the 10th layer. (c) The net cage in the FhSim simulation environment, showing deformation under wave and current excitation. (d) Top view of the fish farm highlighting net cage 11, which accommodates the sensors. A buoy located 400 meters from the net cage measures waves and currents, with the main flow direction indicated. (e) Close-up of net cage 11 showing sensor locations: 5 load shackles measuring mooring line forces and 3 depth sensors measuring net displacement. (f) Side view of the net cage illustrating sensor locations along its depth.

Figure 3: **Multifidelity digital twin for SINTEF ACE fish farm.** (a) The digital twin's workflow is illustrated. The digital twin receives real-time measurements of current speed and direction from a metocean buoy equipped with sensors. These measurements serve as inputs to the low-fidelity Gaussian Process (GP) models, which predict the mooring line loads and the net cage deformation. The low-fidelity predictions, represented by the GP posterior mean values, along with the metocean buoy's current measurements, are subsequently fed into high-fidelity GP models which refine these predictions, achieving greater accuracy that closely matches the actual observations from load shackles and depth sensors. (b) The training stage of the low-fidelity GP models using data from FhSim simulations is described. A standard GP model is trained to map current characteristics to mooring line loads. For net cage deformation, dimension reduction using PCA is performed first due to the extensive data required to describe the flexible net topology. A standard GP model maps current characteristics to PCA coefficients, followed by a reconstruction stage that outputs the predicted net deformation. These low-fidelity GP models are capable of substituting FhSim simulations. (c) The training stage of the high-fidelity GP models based on multifidelity data is illustrated. Limited field sensor measurements randomly selected from January 18 to March 5, 2020, are used for training, including current measurements from the metocean buoy, mooring loads from load shackles, and net cage vertical displacements at depth sensor positions. The current measurements are input into the low-fidelity GP models, and the resulting low-fidelity predictions, combined with the current measurements, serve as inputs for the high-fidelity models. These models are trained to provide predictions that closely match load shackle and depth sensor measurements.

Figure 4: **Multifidelity framework predictions for mooring line loads.** Scatter plots (a), (d), (g) show the low-fidelity GP posterior mean against load shackle measurements. Color scale indicates log count density of data points, with lighter colors representing higher concentrations. The dashed diagonal line signifies perfect prediction accuracy; the red line is the best-fit trend. Low-fidelity GP predictions have broad spread and the red line presents significant deviation from the dashed line, indicating less accurate predictions. Scatter plots (b), (e), (h) show the multifidelity GP model predictions against sensor measurements. Multifidelity GP model has higher concentration of lighter colors near the dashed line, with the red line closely aligning with it, indicating improved accuracy. Plots (c), (f), (i) display predictions over a 36-hour period, excluding the data used for training. Low-fidelity GP posterior mean deviates significantly from actual measurements. Multifidelity GP posterior aligns closely with real observations, capturing the overall trend despite some oscillatory behavior deviations. Mean Absolute Error (MAE) values indicate a significant improvement in prediction accuracy with the multifidelity model compared to the low-fidelity model.

Figure 5: **Multifidelity framework predictions for net cage displacement.** Scatter plots (a), (d), (g) contrast the predictions from the low-fidelity GP model with the depth sensor measurements. The broad dispersion of data points and the notable deviation of the red best-fit line from the dashed diagonal line underscore the model's limited accuracy. Color scale indicates log count density of data points, with lighter colors representing higher concentrations. Scatter plots (b), (e), (h) for the multifidelity GP model. Data points cluster more closely around the dashed line, and the red best-fit line aligns more precisely with the diagonal, indicating enhanced predictive accuracy. Despite the inherent complexities and dynamic motion of the flexible net

cage, the multifidelity model's predictions show improved accuracy. Darker points represent outliers, while lighter points closer to the dashed line signify higher accuracy. The MAE values annotated on these plots quantitatively highlight the multifidelity model's superiority over the low-fidelity model. Plots (c), (f), (i) show the temporal comparison of predicted and actual net cage displacements over a specific period (January 18 to January 27, 2020). The low-fidelity GP model (blue line) exhibits substantial deviations from the actual measurements (green line), particularly for depth sensors 1 and 3. The multifidelity GP model (red line) significantly mitigates these discrepancies, closely mirroring the actual measurements and capturing the underlying trend with greater fidelity.

Figure 6: **Machine-learned surrogate models for predicting net cage deformation.** (a) Comparison of net cage deformation predictions under harsh (301) and mild (801) current excitations by Graph Convolutional Networks (GCNs) with FhSim simulation solutions. (b) Evaluation of net cage deformation prediction accuracy using MAE metric, comparing predictions of each node displacement in the x , y , z directions with the FhSim solution. Two surrogate modeling methods are compared: (i) GP surrogate with prior PCA, and (ii) GCNs. The red line at MAE = 1 serves as a reference point, indicating the threshold beyond which prediction error significantly impacts operations such as UUV control. The surrogates are evaluated across 100 test cases. The final plot summarizes the MAE, showing the average magnitude of errors in predictions for each axis, with error bars representing the variability in the error.