



# Using feedforward neural networks to represent ecosystem dynamics for bioeconomic analysis

Yuanming Ni<sup>1,\*</sup>, Leif K. Sandal<sup>1,2</sup>, Sturla F. Kvamsdal<sup>1</sup>, Cecilie Hansen<sup>3</sup>

<sup>1</sup>SNF – Centre for Applied Research at NHH, 5045 Bergen, Norway

<sup>2</sup>NHH Norwegian School of Economics, 5045 Bergen, Norway

<sup>3</sup>Institute of Marine Research, 5005 Bergen, Norway

**ABSTRACT:** We applied feedforward neural networks to represent ecosystem dynamics that are vital to bioeconomic analysis, ecosystem-based management, or what-if analysis regarding the underlying natural resources. Neural networks are flexible, universal function approximators, recognized for their ability to recover complex nonlinear relationships. In this paper, we treated outputs from an end-to-end Atlantis model as synthetic data and used them as training data for the neural networks. After learning the seasonal dynamics of a multispecies system, we forecasted system states with different sets of specified harvest policies using the trained networks. We demonstrate that neural networks can capture key dynamics in part of the ecosystem efficiently, and give fast updates of states that are needed for optimization and decision-making. The trained networks are reduced and more flexible systems compared to the large-scale simulator model, which is more costly to run and does not have a format allowing human actions in the form of feedback policies, or harvest control rules, i.e. decisions depending on states as well as time.

**KEY WORDS:** Neural networks · Ecosystem dynamics · Bioeconomic · Seasonality · Ecosystem-based management · Cod

## 1. INTRODUCTION

Ecological systems have inherent nonlinear and complex dynamics (Burkett et al. 2005) and exact, closed-form representations of the dynamics may be fundamentally unknown (Holt et al. 2014, Planque 2016). Many natural resource industries depend on ecosystem functioning, making models or representations of ecosystem dynamics necessary for bioeconomic studies. However, these representations often rely on assuming dynamic functional forms and statistical regression analysis with limited observations, which may not capture key ecosystem complexities. Representations in terms of explicit equations are intuitive, intellectually stimulating, and align with time-honored modeling approaches (May et al. 1979), but require considerable efforts to deal with issues

such as parameterization and model fit (Kvamsdal & Sandal 2015, Ekerhovd & Kvamsdal 2017). Moreover, the assumptions and simplifications inherent in closed-form models may lead to less relevant management decisions.

Neural networks (NNs) are recognized for their prediction power and ability to infer hidden patterns and behaviors (Joseph 2020). They are adaptive and flexible nonlinear mapping structures that are 'universal approximators' for any data, which is especially practical when the underlying relationships are unknown (Lek & Guégan 1999). By using NNs fitted to stock dynamics, we circumvent the issues related to imposed dynamic equations. Properly trained NNs can surmount some of the fundamental limitations of closed-form models but still maintain attributes required for bioeconomic analysis, such as flexibility and speed.

\*Corresponding author: [yuanming.ni@snf.no](mailto:yuanming.ni@snf.no)

Recognized as powerful instruments for ecological modeling, machine-learning methods are increasingly pursued for their potential in advancing ecosystem studies (Willcock et al. 2018, Malde et al. 2020, Scowen et al. 2021). Applications of NNs in ecological modeling are widespread, including emulation of climate models (Krasnopolsky et al. 2005), prediction of population outbreaks or species distribution (Rammer & Seidl 2019, Deneu et al. 2022), simulation of various types of ecosystem components (Lu et al. 2016), and evaluation of fisheries management (Russo et al. 2014, 2019). However, most of these applications lend themselves to what-if type analyses where the research objective is scenario simulation or system prediction.

Our present objective, the representation of ecosystem dynamics in bioeconomic models, requires us to take into consideration several features beyond the skill of capturing dynamic behavior. This skill is crucial for the relevance of a model, but we need to further consider dimensionality and run time. The challenges brought by the curse of dimensionality and computation time may be waning as advanced optimization methods are being developed (Cai & Lontzek 2019, Springborn & Faig 2019), but we are still limited by them. Bioeconomic models are useful for decision analysis, incorporating links between biological and economic production, but are in most settings only feasible with a somewhat limited dimensionality, that is, involving a handful of dynamic variables (Crépin et al. 2011, Levin et al. 2013, Kvamsdal & Sandal 2015). Run time is essential on account of the iterative procedures used in decision analysis and related work (Ni & Sandal 2019, Kvamsdal et al. 2020b). Recent techniques, for example reinforcement learning (Malo et al. 2021, Tahvonen et al. 2022), involve decision-making algorithms that are iterative as well.

A comprehensive empirical basis is essential for NNs to be effective, making data scarcity a fundamental challenge, particularly for marine ecosystems that are difficult to observe. The limitation of empirical evidence is severe for extensive ecosystem models, as well as for compact models (Ekerhovd & Kvamsdal 2017). The lack of observations is further aggravated by the non-ergodicity of ecosystems, that past states cover only a small fraction of possible future states (Planque 2016). Nevertheless, comprehensive end-to-end ecosystem models, such as the Atlantis model (Audzijonyte et al. 2019, Fulton et al. 2011), exist and can provide a synthesized reality that is relevant for ecosystem-based management. We acknowledge the skill of state-of-the-art Atlantis models

to capture the general nature of ecosystem dynamics (Olsen et al. 2016). These large-scale models are convenient vehicles for summarizing and communicating data and knowledge in a form that can be captured by models such as NNs. We refer to the outputs from Atlantis models as synthetic data. Rather than using limited empirical data, we use relatively abundant synthetic data to train NNs and thereby also achieve better control for input data quality.

We aim to establish mathematical structures, in the form of NNs, that can represent the dynamics in the training data to a sufficient degree, such that key complexities are captured. Meanwhile, the representation should be compatible and computationally efficient to allow for bioeconomic analysis, feedback policies, optimization, or what-if analysis. The high level of complexity and dimensionality of models such as Atlantis models makes applications in such analyses infeasible; the levels of aggregation are conflicting. NNs can readily be defined on the appropriate aggregation level. The proposed NN approach could enlarge our toolbox when efficient updates of a system are needed.

## 2. MATERIALS AND METHODS

### 2.1. Earlier work on neural networks for ecosystems

The main effort thus far to apply NNs fitted to marine ecosystem dynamics, with particular application in bioeconomics, has been that of Russo et al. (2014). A similar but further developed approach was used in Russo et al. (2019). The NNs in these studies represent development of ecosystem components (age classes of fish stocks, harvest effort, and environmental conditions) at a relatively detailed spatial scale. For example, the spatial grid used in Russo et al. (2014) contains almost 400 cells. The level of detail and aggregation places their NNs closer to the Atlantis model than what we presently set forth to develop. As shown by Russo et al. (2014, 2019), the networks support what-if analysis of changes in management, for example with respect to spatial closures. Our aim, however, is primarily to develop networks to support bioeconomic analysis and other applications where decision-making is needed, particularly under uncertainty. Our example networks will be defined on an aggregation level relevant for this purpose.

The level of aggregation is a key feature for ecosystem-based management problems, and thus for

the formulation of relevant models. Consider, for example, the use of harvest control rules (Kvamsdal et al. 2016). A harvest control rule is typically formulated on the stock level and has a feedback nature; it provides the level of harvest as a function of the stock level. A decision model for analysis of harvest control rules will then naturally be formulated with stock levels as state variables. In comparison, a stock in the NNs formulated by Russo et al. (2014), as well as in Atlantis models (Hansen et al. 2019b), is represented by several age classes distributed across a spatial grid. Thus, many different states in the Russo or Atlantis model corresponds to a given point in the state space of the decision model. The lack of correspondence between state spaces is a problem for what-if analysis of harvest control rules and optimization methods. That is, the mismatch in dimensionality makes direct application of highly detailed models infeasible in decision-making frameworks. Put differently, the ecosystem model of the subsystem must be designed to fit the decision framework. While the models developed by Russo et al. (2014, 2019) serve a similar purpose to ours, they differ in their compatibility with bioeconomic analysis methods such as stochastic optimization (Sandal & Steinshamn 1997).

## 2.2. Synthetic data

The synthetic data used for training feedforward NNs (i.e. NNs where the information flows in only one direction; that is, forward) are provided by the Nordic and Barents Seas Atlantis model (NoBa Atlantis), which includes 53 species and functional groups, a grid of 60 polygons covering an area of roughly 4 million square kilometers, and multiple layers representing the environment and its ecosystem in the Nordic and Barents Seas (Hansen et al. 2019b). NoBa Atlantis is an end-to-end ecosystem model that simulates daily changes in the system and is forced by daily input of temperature, salinity, and volume fluxes from a set of regional ocean modeling system (ROMS) models (Shchepetkin & McWilliams 2005). It uses an RCP 4.5 scenario for the climate projections (Sandø et al. 2014, Skogen et al. 2018), which can be categorized as optimistic (Hansen et al. 2019b), with a temperature increase in the Barents Sea at around 1°C over the simulation period (i.e. by 2070). The ecosystem components are linked together through a diet matrix, where the overlap between the predator and prey is defined. Harvesting is applied at a constant rate within a year and to the entire stock, including juveniles, for all commer-

cial species. NoBa Atlantis updates daily with some processes even at smaller internal time steps, and the outputs are printed and stored 5 times per year.

A simple harvest approach has been applied to the simulations in Hansen et al. (2019b), which are used as data for training the NNs. For the historical simulation period (1981–2016), the historical harvest rates have been calculated from stock assessments (ICES 2020, 2018) for the commercial stocks from 1981 to 2016 and are applied in Atlantis simulations without modifications. For the future simulation period (2017–2068), the maximum sustainable yields (MSY) and the corresponding fishing mortality rates ( $M_{MSY}$ ) are calculated for the commercial species within NoBa Atlantis (Hansen et al. 2019b). The  $M_{MSY}$  values are multiplied with 4 scalars—0.6, 0.8, 1.0, and 1.1—forming 4 different harvesting scenarios throughout the future simulation period. For the scalar of 1.0, fishing occurs at MSY levels for all harvested species; for a smaller scalar, for example 0.6, fishing pressure is reduced, and the scenario reflects sustainability considerations. For every scenario, there are 14 runs, each following an individual pattern of mesozooplankton growth. This introduces some degree of randomness in the simulated results under the same harvesting scenario. NoBa Atlantis has been extensively calibrated and tested, with respect to diets, population development, individual weights, total biomass, catches, patterns, and distributions, for those species where available observations could be found (Hansen et al. 2019a,b)

The key commercial fish in the Barents Sea is the Northeast Arctic cod, around which several related species are often modeled together in bioeconomic analysis (Helstad 2001, Aanestad et al. 2007). In this study, we employed feedforward NNs to learn the dynamics of a 3-species subsystem consisting of cod, capelin, and juvenile Norwegian spring-spawning herring (hereafter herring). Cod is the predator of the other two, and there is also predation on capelin larvae by juvenile herring. Note that our biomass model specifies the juvenile part of the herring stock because the juveniles interact with cod and capelin in the Barents Sea, while the adult herring are mostly found in the Norwegian Sea. From here on, juvenile herring refers to one state in the 3-dimensional system. This setup corresponds to the models in Durant et al. (2008), Kvamsdal & Sandal (2015), and Kvamsdal et al. (2020b). While the trained networks are modeled to represent the dynamics of the subsystem, the training data from NoBa Atlantis reflect the full ecosystem reality and contain comprehensive ecosystem knowledge.

### 2.3. Training of feedforward neural networks

Feedforward NNs feature a straightforward architecture without cycles or loops. A standard feedforward NN comprises 3 layers: an input layer responsible for receiving data; a hidden layer that incorporates nonlinearity through a nonlinear activation function; and an output layer that applies a linear activation function. NNs estimate internal parameters, called weights and biases. Compared to the estimates from traditional regression equations, these internal estimates are less intuitive to interpret. Hyperparameters are external parameters that determine the network structure and the learning process. To ensure successful training, key hyperparameters need to be tuned (Scowen et al. 2021) according to the specific problem, for example, through grid search or trial and error (Hutter et al. 2015).

We built a feedforward NN structure with 3 layers using ‘fitnet’ in the MATLAB neural network fitting tool. We adopted the Rectified Linear Unit (ReLU) activation function (Goodfellow et al. 2016) for the hidden layer. It is a piecewise linear function that has emerged as the default choice for many types of NNs due to its ease of training and superior performance to traditional sigmoid activation functions (Javid et al. 2020). We applied a process function of ‘mapstd’ that normalizes the inputs and targets such that they have zero mean and unit standard deviation. The performance function measures the accuracy of network prediction, with mean squared error (MSE) serving as the default metric. A smaller MSE corresponds to a higher degree of accuracy in reproducing the training examples. In our problem, MSE is applied in conjunction with standard normalization, which normalizes errors between  $-2$  and  $2$ , a feature suited for networks with multi-element outputs. The training algorithm minimizes the normalized MSE and the default algorithm for ‘fitnet’ is Levenberg–Marquardt backpropagation (Møller 1993).

One critical matter in NN modeling is mitigating overfitting to ensure that the trained network can generalize well to new, previously unseen data (Goodfellow et al. 2016). Weights within a NN specify the strength of connections between nodes and convergence to large weights is often a sign of overfitting, implying significant importance of specific input examples that is likely to impede network generalization. To evaluate the networks’ ability to generalize, we adopted the holdout method, that is to withhold 10% of the data as an independent test set, untouched during training. We utilized MSE for the holdout test set and the sum of squared weights as

metrics for assessing network generalization, with lower values signifying lower risk of overfitting. Furthermore, under each network specification, we retrained 5 networks starting with different parameter initializations and initial data divisions. The final simulation outcome for a given specification was derived by averaging the predictions of the 5 networks.

Ecosystem data, available in the form of long time series, were treated as iterative mappings of the system, providing input examples for the transition dynamics. While we acknowledge critique concerning the reliability of long-term ecosystem projections, short-term predictions can be accurate (Planque 2016, p. 205). And the credibility of our model only relies on the quality of one-time-step transitions because we treated synthetic time series data as successive realizations of the same data-generating process over a time step.

Given 5 observations per year in the data, we considered there to be 5 seasons, noted by a season indicator  $s$ . The total annual harvest was determined at the beginning of a year by multiplying the stock biomass with an annual  $M_{\text{MSY}}$ . One-fifth of this determined amount is removed from the stock each season, that is, biomass removal is constant throughout the year. We assumed that system transitions can be captured by the following general relationship:

$$\mathbf{x}_t = F(\mathbf{x}_{t-1}, s) - \mathbf{h}_t \quad (1)$$

where  $\mathbf{x}_t$  is the biomass vector of the system at time  $t$ ,  $\mathbf{h}_t$  is the harvest vector of the system during  $(t-1)$  to  $t$ , and  $F$  is the transition function of the system. The mapping function  $F$  can be approximated by:

$$F(\mathbf{x}_t, s) = \mathbf{x}_{t+1} + \mathbf{h}_{t+1} \quad (2)$$

To achieve the approximation, the vector of the biomass  $\mathbf{x}_t$  and the corresponding season indicator  $s$  are fed into NNs as training inputs; total biomass before harvest at the next time step,  $(\mathbf{x}_{t+1} + \mathbf{h}_{t+1})$ , and the season indicator of  $(t+1)$  are set as the training targets. In total, there are 17 010 examples of system transitions: 2506 (36 yr, 5 observations per year, 14 runs) during the historical simulation period and 14 504 (52 yr, 5 observations per year, 14 runs, 4 harvest scenarios) during the future simulation period. The seasonal effects in the dynamics are picked up by the NNs automatically through the season indicator  $s$  in the learning process.

Fig. 1 provides a detailed illustration of our network structure. The input layer has 4 nodes: 1 each for cod, capelin, and juvenile herring biomass levels, and 1 for season. The output layer has 3 nodes repre-

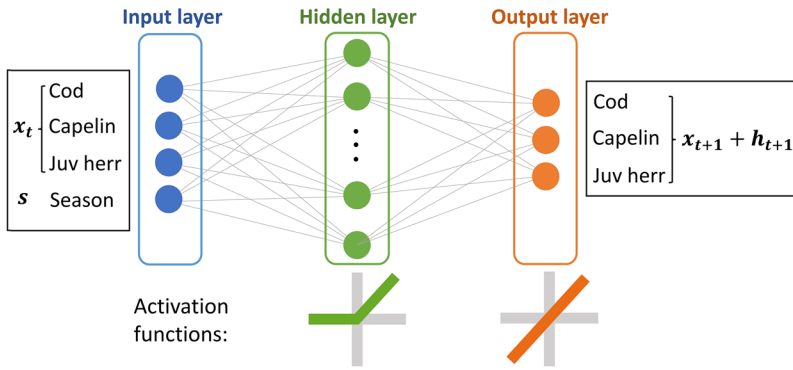


Fig. 1. Neural network (NN) structure: 4 nodes in the input layer representing biomasses for the three species and the season indicator, a varying number of nodes in the hidden layer, 3 nodes in the output layer representing sums of biomass and harvest for the three species. Thin grey lines: connection weights between nodes. Rectified linear unit (ReLU) activation function in the hidden layer, and linear activation function in the output layer (diagrams at bottom)

senting the biomass levels before harvesting in the next period. The number of hidden nodes, a key hyperparameter not specified in Fig. 1, dictates the structural complexity and overall predictive power of the network. According to the universal approximation theorem, a feedforward NN with only one hidden

layer is capable of approximating any relation given enough hidden nodes (Hornik et al. 1989). However, overfitting may occur if the network’s capacity significantly exceeds the demands of the problem at hand. Thereby, we implemented a grid search to optimize the number of hidden nodes, striking a balance between prediction accuracy and generalization, with the results shown in Fig. 2.

Fig. 2a reveals that the MSE for the holdout test set is higher than that of the training set, indicating larger errors when predicting outside the learned examples. The MSE for the training set decreases as the number of hidden nodes increases, but this trend plateaus beyond 30 nodes. Fig. 2b shows that both the MSE for the holdout test set and the sum of squared weights decrease with increased network complexity until around 30 hidden nodes. This is partly due to a mechanism in the training algorithm called early

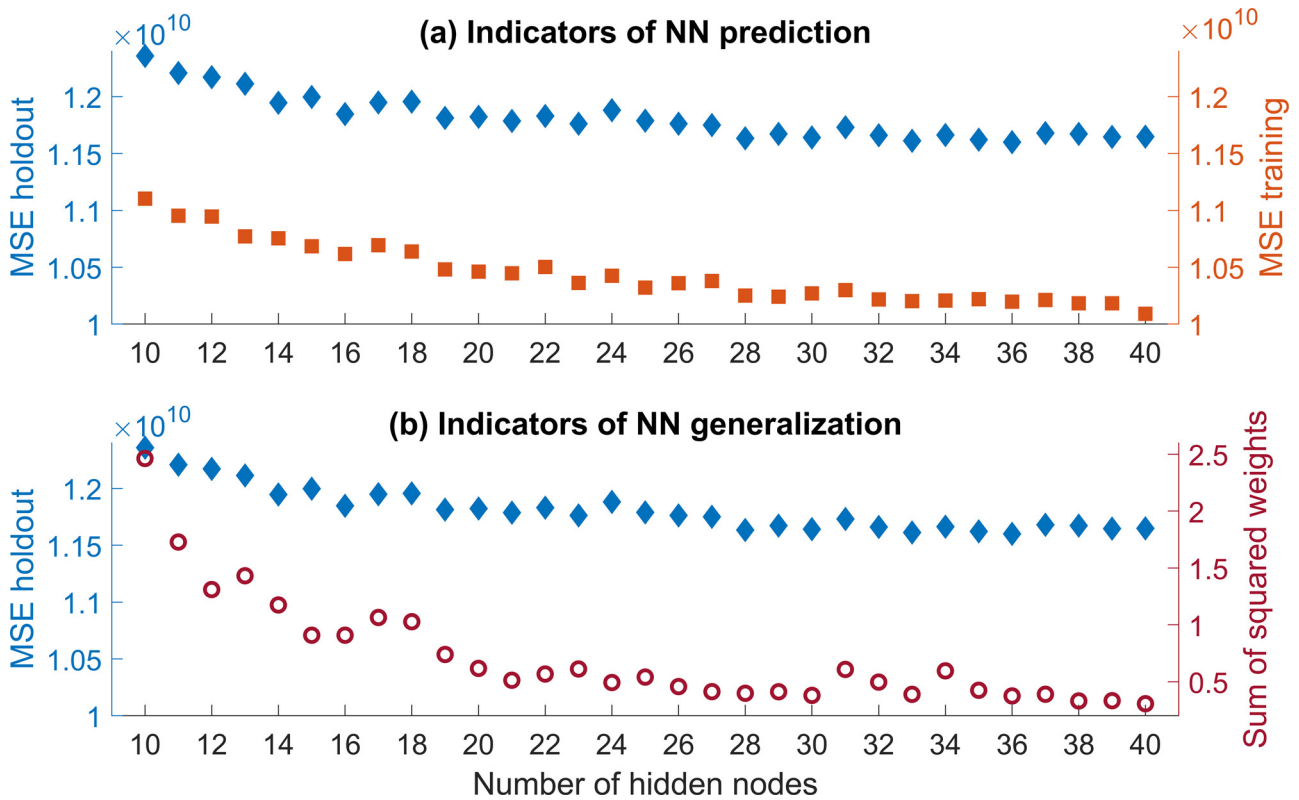


Fig. 2. Measures of neural network (NN) performance. (a) Indicators of NN prediction (orange squares: maximum sustainable yield, MSE, of the training set; blue diamonds: MSE of the holdout test set). (b) Indicators of NN generalization (blue diamonds: MSE of the holdout test set; red circles: sum of squared weights). Note that for better visualization of the results, the MSE is not normalized in this figure

stopping that prevents overtraining and overfitting with a large number of nodes. These observations suggest that beyond 30 hidden nodes, marginal improvements in prediction accuracy are observed without further gains in network generalization. Consequently, in the following applications, we selected 3 network structures with 35, 36, and 37 hidden nodes respectively and used the average of the 15 outputs (note there are 5 networks trained per structure) as the final output. The absence of abnormally large weights reinforces our argument that the selected networks are trained properly and can be used for simulation and forecasting.

#### 2.4. Neural network forecasting

Prediction using NNs is referred to as forward propagation or simulation, which is usually straightforward and fast (Krasnopolsky et al. 2005). Iterative simulations in a feedback loop over a longer period is what we call NN forecasting. To be more specific, NN forecasting is achieved by calling the trained NNs repeatedly in a closed loop, using outputs from the previous simulation as basis for inputs for the next simulation. Since the output layer for regression problems has a linear activation function, extreme or meaningless outputs could emerge, especially when NNs extrapolate outside of the data range. The likelihood of having paths leaving the data range increases as the prediction errors accumulate over time. Thus, we applied a lower boundary to prevent negative values for the new state after harvesting. Any negative new state was set to a small number (2000 tonnes) to avoid an absorbing zero boundary.

As uncertainties are pervasive in fisheries and ecosystem dynamics, incorporating stochasticity can enhance our analysis to be more relevant for real-world situations. We applied stochastic NN forecasting by adding a multiplicative noise factor to the process:

$$\mathbf{x}_{t+1} = \mu F(\mathbf{x}_t, t) - \mathbf{h}_{t+1} \quad (3)$$

where the noise factor  $\mu$  is a 3-dimensional vector with independent components, each drawn from a truncated normal distribution such that  $0.9 < \mu < 1.1$ . Put differently, the networks' final output will, at each time step, experience a random multiplicative kick, which could be positive or negative, of a maximum of 10% in scale during forecasting. We advocate for imposing kicks that are a maximum of 10% based on its alignment with the ICES stock assessment (ICES 2021). In the assessment of cod, the estimated spawning stock biomass (SSB) is presented in

conjunction with a high value and a low value, which are typically 10 to 20% deviated from the mean. Therefore, the formulation in Eq. (3) aligns with established practices, which bolsters relevance and reliability of the NN forecasting outcomes.

### 3. RESULTS

In this section, we present 3 distinct analyses: (1) the application of NN forecasting using harvest rates derived from the training data, incorporating both deterministic and stochastic approaches; (2) the implementation of stochastic NN forecasting with state-dependent, adaptive harvest rates, updated on an annual or seasonal basis; and (3) reintroducing the yearly updated harvest rates from the second analysis back into the Atlantis model and comparing the simulation outcomes between the 2 representations. Analogous to the first analysis, which compares the 2 models using the same set of harvest rates, the innovative aspect of the last analysis is that the realized harvest rates are not sourced from the training data. Instead, they comprise previously unseen, new policies both for the NNs and the Atlantis model.

In analyses 1 and 2 concerning stochastic NN forecasting, we simulated 100 trajectories for each policy specification and calculated the corresponding mean path and standard deviations. In the first analysis, the stochastic forecasting process runs for the historical and future simulation periods separately: the mean state of the 100 runs at the end of the historical period is used as the initial state for the future simulation period. Note that the first state in 1981 is the only external input to initiate the simulations. The NN system is run for 440 time steps, free from external calibrations. In the second analysis, forecasting is applied in one go for the entire time span, with updating of harvest rates occurring at different frequencies, yearly and seasonally.

#### 3.1. Forecasting with trained neural networks

The primary objective of the first analysis is to evaluate whether the system trajectories forecasted by NNs, given the same harvest rates, resemble the original data paths. We applied harvest rates from the training data and we select one out of 4 harvesting scenarios as in Hansen et al. (2019b). The implemented harvest rates in this analysis are historical harvest rates during the historical simulation period

(1981–2016) followed by a constant harvest rate of  $0.8 \times M_{MSY}$  during the future simulation period (2017–2068).

Fig. 3 displays the forecasted progression of the subsystem in the absence of stochasticity, achieved by excluding the noise parameter  $\mu$  from Eq. (3). The forecasted paths demonstrate that NNs can replicate, to a large degree, the data trajectories of cod, capelin, and juvenile herring under the same harvest policy. For the cod stock, the forecasted path largely resembles the observed data in terms of behavior, with regard to both seasonal variations and interannual fluctuations. The high level of synchronization between network outputs and data paths is more evident in the historical simulation period. For the future simulation period, the forecasted path is mostly at a constant level with regular seasonal variations.

Concomitantly, Fig. 3 includes a 2-SD range surrounding the mean, derived from the 14 runs that constitute the dataset. Notably, this interval is particularly narrow for cod. The narrow range of data means that cod is fairly robust to the variations in mesozooplankton forcing that the NoBa Atlantis model incorporates. However, cod is under the influ-

ence of additional stochastic forces and a closer fit of the NN forecast to the data would probably be overfitting. Capelin, as a pelagic stock, has high variability in general, and the forecasted path could capture some of the dynamics in the historical simulation period but less in the future simulation period. It seems that given a constant harvest rate, the capelin stock simulated by the NNs has stable seasonal cycles. For juvenile herring, the range of data covers a wide band which includes the NN path most of the time. The biomass of juvenile herring keeps shrinking during the entire period, and the seasonal variations are reduced and smoothed towards the end when stock is collapsed, which is reflected in the forecasted path.

Fig. 4 demonstrates forecasted paths with stochasticity, following the model in Eq. (3). We observed that the mean paths of the 100 runs depicted in Fig. 4 are essentially equivalent to the paths in Fig. 3, indicating that the average stochastic trajectories are highly representative of their deterministic counterparts. The 2-SD bands of the simulated paths cover the range of data paths completely for the cod stock. For capelin, some peaks and troughs of seasonal

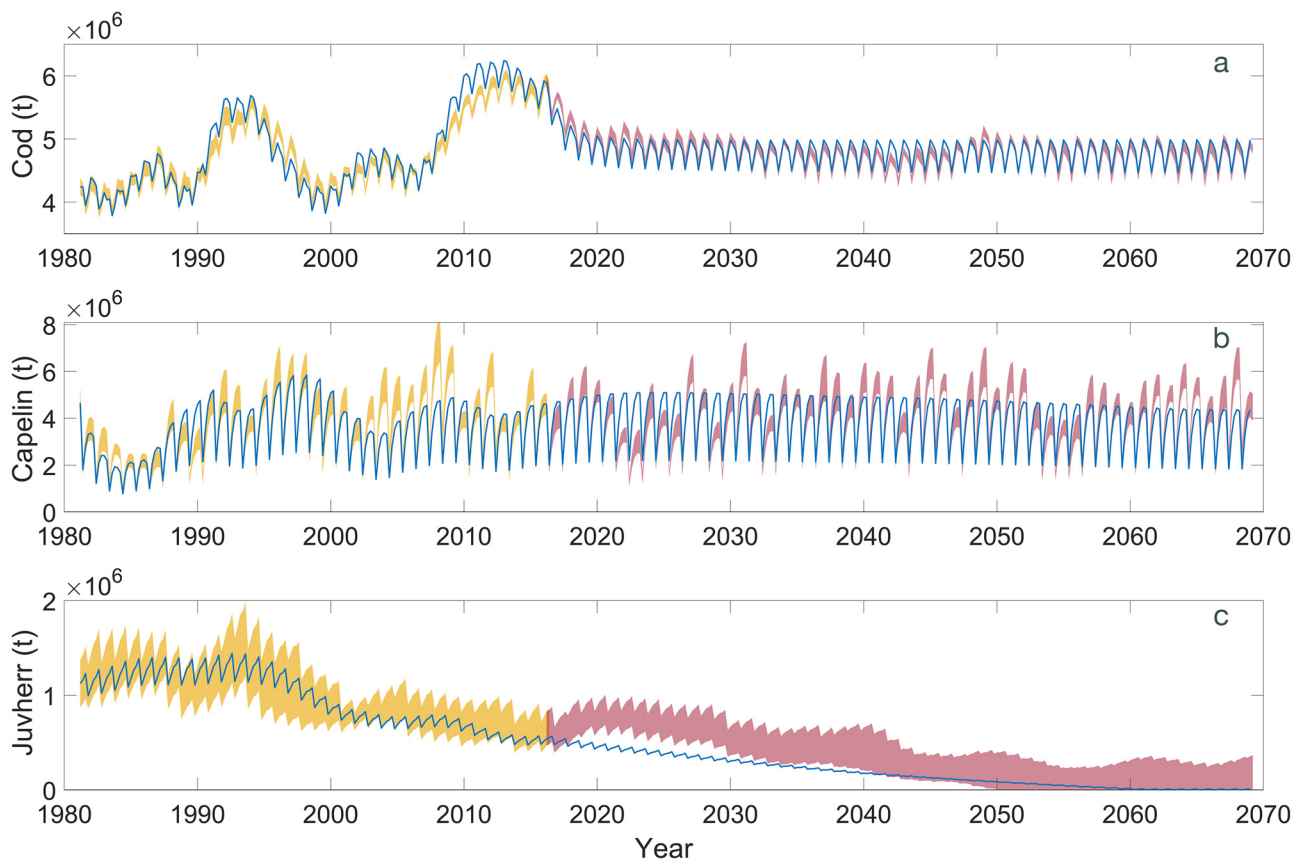


Fig. 3. Paths created by NN forecasting (blue solid line) plotted against the 2-SD range of data (yellow band: historical simulation period 1981–2016; red band: future simulation period 2017–2068) for (a) cod, (b) capelin, and (c) juvenile herring

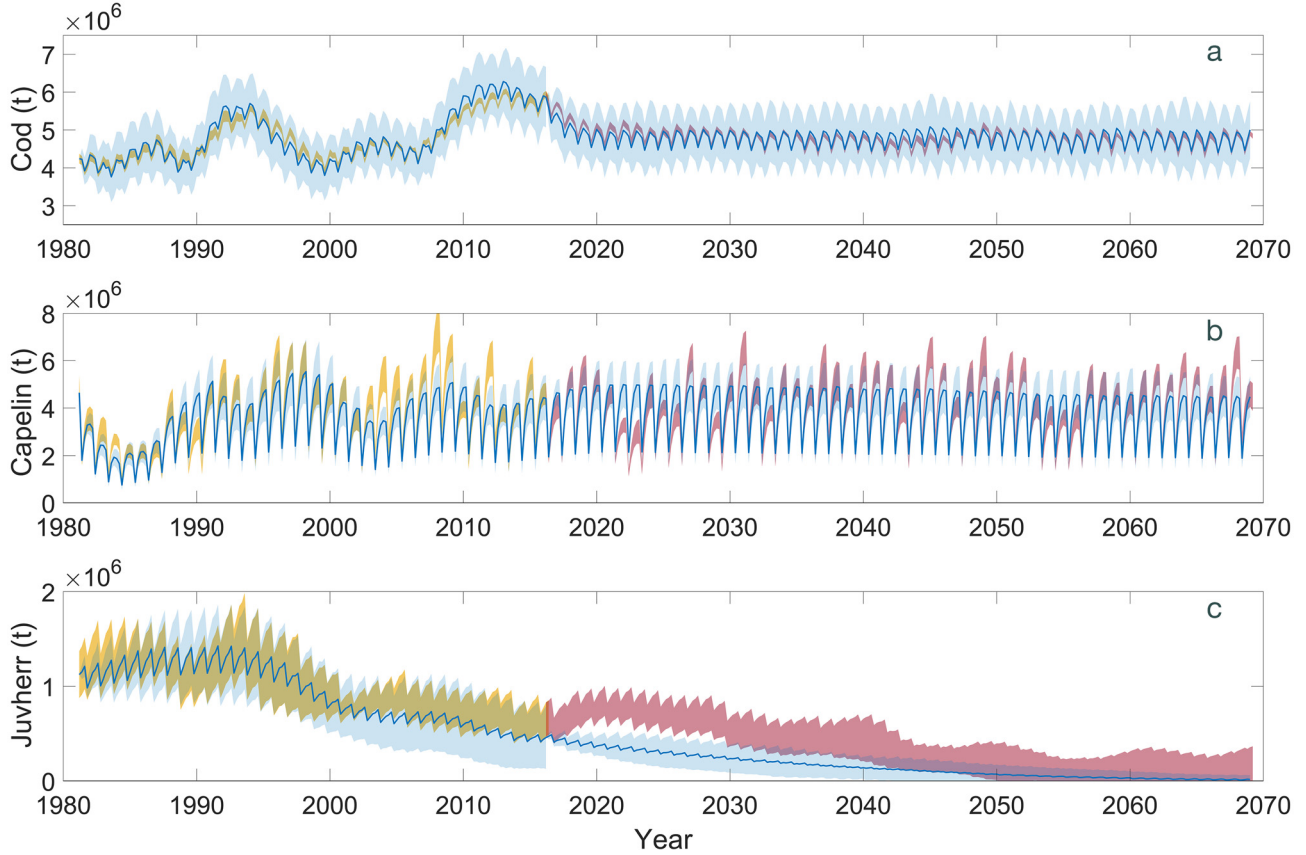


Fig. 4. Mean paths created by NN forecasting (blue solid line) plotted against the 2-SD range of 100 simulated paths (light blue band) and the 2-SD range of data (yellow band: historical simulation period 1981–2016; red band: future simulation period 2017–2068) for (a) cod, (b) capelin, and (c) juvenile herring

variations in the data stretch beyond the scope of the simulated paths bands. It seems that the NNs generate seasonal dynamics of capelin that have a stronger contraction toward an equilibrium level compared with the Atlantis model. During the historical simulation period of juvenile herring, the 2-SD bands of the data and forecasted paths overlap to a large extent, but the forecasted values seem to be dragged down during the future simulation period. This downward dragging phenomenon is possible when stochasticity is present, especially when the noise term is multiplicative (Poudel et al. 2015).

### 3.2. Forecasting under an alternative management policy

To demonstrate that the trained NNs can be applied for relevant decision analysis, we present a second analysis where we specify harvesting rules with adaptive features (Ussif & Sumaila 2005). We applied a simple feedback policy without interaction

among species; that is, the harvest decision is independent of information on other species. For a given stock, the updated harvest  $h_{\tau+1}$  is adapting according to the state  $x_{\tau}$ , the target biomass level  $B_{\tau}$  and the harvest  $h_{\tau}$  in the previous period, as well as a contraction parameter  $\eta$ , as given in Eq. (4):

$$\begin{cases} h_{\tau+1} = \max(0, h_{\tau}(1+z_{\tau})) \\ z_{\tau} = \eta \left( \frac{x_{\tau}}{B_{\tau}} - 1 \right), \quad \eta > 1 \end{cases} \quad (4)$$

Continuous implementation of this harvest rule will contract the stock toward the target biomass level at a speed determined by  $\eta$ : a larger  $\eta$  indicates a faster contraction.

We consider both seasonal and yearly updating of  $h_{\tau+1}$ . While yearly decisions over the harvest level are in line with current fisheries management practice, there may be good reasons for considering intra-annual management interventions (Smith 2012, Huang & Smith 2014). When the harvest is updated seasonally, the unit of  $\tau$  is season; when updating happens yearly, the unit of  $\tau$  is years. Similar to the



simulations in the Atlantis model, yearly updated harvest implies a constant biomass removal throughout the year, based on the predetermined amount at the beginning of the year. Conversely, seasonal updating entails adjusting harvest levels more frequently, 5 times a year, in response to the situation in the previous season.

Upon each updating of harvest, we encounter 1 of the 3 situations, as summarized in Eq. (5):

$$\left\{ \begin{array}{l} \frac{x_\tau}{B_\tau} \geq 1 \quad \Leftrightarrow z_\tau \geq 0 \Leftrightarrow h_{\tau+1} \geq h_\tau \\ \frac{x_\tau}{B_\tau} \leq \left(1 - \frac{1}{\eta}\right) \Leftrightarrow z_\tau \leq -1 \Leftrightarrow h_{\tau+1} = 0 \\ \left(1 - \frac{1}{\eta}\right) < \frac{x_\tau}{B_\tau} < 1 \quad \Leftrightarrow -1 < z_\tau < 0 \Leftrightarrow 0 < h_{\tau+1} < h_\tau \end{array} \right. \quad (5)$$

When the state-target ratio  $\left(\frac{x_\tau}{B_\tau}\right)$  is larger than 1, the updated harvest  $h_{\tau+1}$  will be higher than harvest in the previous period,  $h_\tau$ ; when the state-target ratio is lower than a critical moratorium level of  $\left(1 - \frac{1}{\eta}\right)$ , the updated harvest is zero; and when the state-target ratio is between the moratorium level and 1, the updated harvest is positive but smaller than the harvest in the previous period. A stronger contraction (larger  $\eta$ ) implies a higher precautionary stock level to activate the moratorium. In this way, the harvest rule represented in Eq. (4) will guide the system toward the target biomass levels in a feedback fashion. Note that  $h_\tau$  as harvest in the previous period refers to the most recent non-negative harvest such that zero harvest, which is activated once the stock falls below the moratorium level, will not be carried forward indefinitely even when the stock has grown back.

The harvest control rule in Eq. (4) bears resemblance to the current fisheries management in the Barents Sea (Kvamsdal et al. 2016). The prevailing management plan for cod in the Barents Sea (ICES 2021) is determined by comparing the current SSB to a precautionary reference SSB level. When the observed SSB is at or above the precautionary level, the exploitation rate is increased, and conversely, it is reduced when the SSB falls below the reference.

In this analysis, initial biomass levels are proximate to Atlantis model states in 1981 and initial harvests are calculated using the  $M_{MSY}$  (0.4 for cod and 0.05 for capelin) from Hansen et al. (2019b). We applied 2 different contraction speeds for cod:  $\eta = 1.2$  and  $\eta = 3$ . used a small contraction parameter  $\eta = 0.1$  for the capelin stock. There is no harvest of juvenile herring in this analysis, following real-world practices. In Fig. 5, we display the stochastic paths for cod and capelin forecasted by NNs and the corresponding cod harvest under the 2 different contraction speeds

and updating frequencies. The harvest of capelin, being negligible in both volume and value, is not reported.

For the cod stock biomass, we observe from Fig. 5a,c that with more frequent updating of harvest, the state contracts toward the target faster and leads to more frequent oscillations around the target. Towards the end of the simulated period, the biomass trajectories of both seasonally and yearly updated policies end up in the vicinity of the target level. For cod harvests in Fig. 5b,d, the average trajectories appear smoother with the yearly updated policy, and exhibit fewer fluctuations when the contraction is slow. Some delay between harvest and biomass development is expected, and it is more noticeable when the contraction speed is slower and when policy updates are less frequent.

One interesting observation regarding the cod stock is that the fast contraction in the harvest rules has a similar effect as strong stochasticity in the state simulations, especially when updating occurs frequently. The 2-SD band of the seasonally updated policy under fast contraction is the widest throughout the 15 yr period we considered and even falls below the moratorium level several times. The effects of strong contraction in the policy are twofold: significant reduction in harvest when the state is low and aggressive harvesting when the state is high. While the former is less of a concern from a conservation aspect, the latter, together with stochasticity, could lead to a dangerous combination of high fishing pressure and overestimated stock level or abruptly reduced stock biomass. Therefore, in the presence of stochasticity, aggressive harvesting should be dealt with cautiously, even when the stock seems abundant above the target level.

The aforementioned observations and interpretations are sensible and intuitive, demonstrating that the trained NNs are capable of simulating relevant system dynamics, which could potentially inform policymaking. However, this analysis is only one instance of a specific type of functional form representing the policy. It is tractable to assess a broad range of candidate policies within a reasonable timeframe using NNs, thereby laying the foundation for relevant bioeconomic analysis.

### 3.3. NoBa Atlantis simulations under the alternative management policy

To further test whether the trained NNs can capture the dynamics embedded in NoBa Atlantis out-

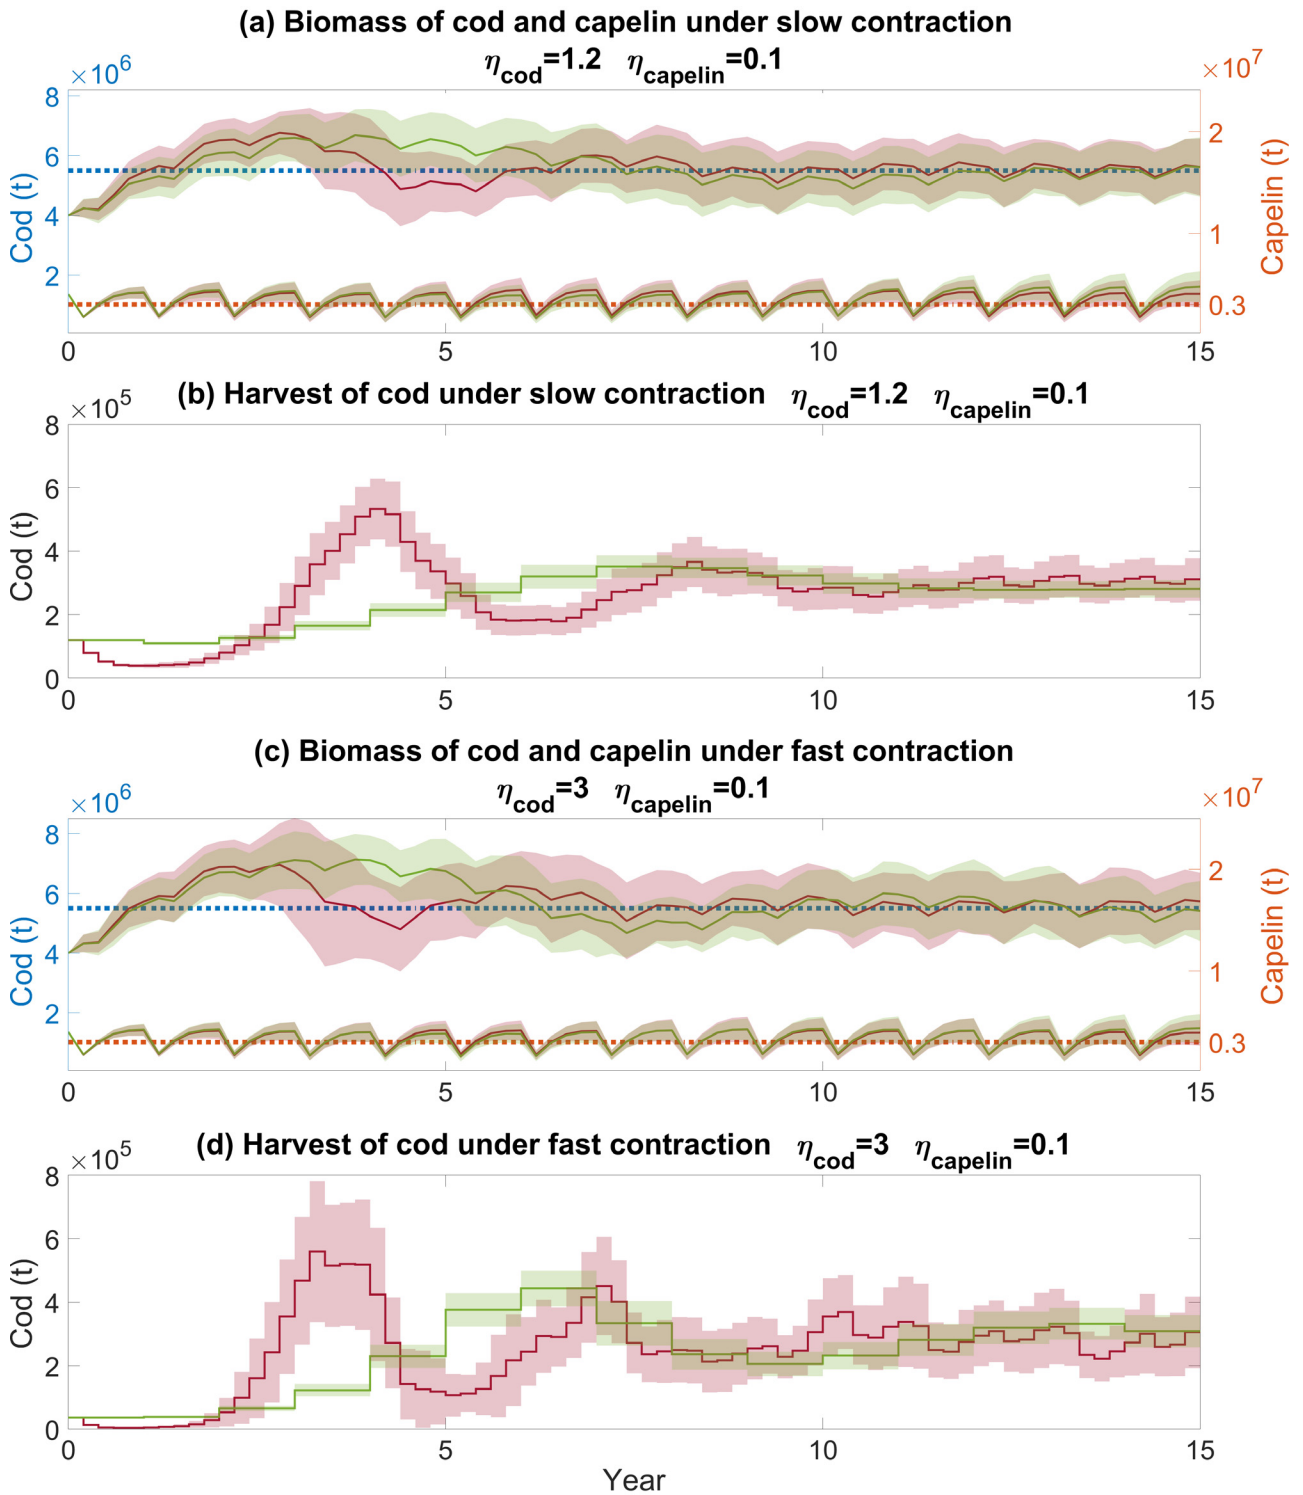


Fig. 5. (a,c) Cod and capelin biomass level paths under stochastic NN forecasting and (b,d) the corresponding cod harvest given 2 contraction speeds  $\eta$ . Red solid line: mean path of biomass or harvest under the seasonally updated policy. Green solid line: mean path of biomass or harvest under the yearly updated policy. Light red band in a and c: 2-SD of the 100 biomass paths under the seasonally updated policy. Light green band in a and c: 2-SD of the 100 biomass paths under the yearly updated policy. Light red band in b and d: 0.5-SD of the 100 harvest paths under the seasonally updated policy. Light green band in b and d: 0.5-SD of the 100 harvest paths under the yearly updated policy. Blue dashed line: target biomass level for cod. Orange dashed line: target biomass level for capelin

puts, we recorded the harvest rates realized during the NN forecasting process of the previous analysis (section 3.2). We fed these harvest rates back into the Atlantis model and compared the simulation outcomes. This particular version of NoBa Atlantis does not have a harvesting module that can deal with feedback policies during an ongoing simulation. Consequently, the harvest rates are not adaptive within Atlantis; however, when the harvest rates are equal in both models, they provide a robust foundation for comparing the Atlantis model with the NNs representation. In this analysis, only yearly updated harvest rates are utilized (see green lines and bands in Fig. 5), as Atlantis enforces constant biomass removal within a year. For commercial species other than cod and capelin, harvesting at  $M_{MSY}$  is implemented in Atlantis throughout the 15 yr period. Fig. 6 illustrates the forecasted biomass trajectories of cod and capelin by both models.

The average biomass trajectories of cod overlap to a large extent between the 2 models: both exhibit highly comparable seasonal variations and inter-annual trends. These observations serve as compelling evidence that, when provided with the same harvest rates, trained NNs can produce cod paths nearly identical to those of the Atlantis model. The cod paths created by Atlantis are slightly smoother, exhibiting lower peaks and shallower troughs compared to those generated by NNs. A potential inference of this phenomenon is that NNs represent a reduced, lighter system that is more dynamic, while Atlantis simulations carry forward extensive details and dimensions, leading to an averaging or smoothing effect when aggregated to the stock level.

Within a range of 2 SD, the NN predictions of cod biomass completely encompass the Atlantis projections. This observation highlights the differing manifestation of randomness in the 2 models: Atlantis accounts for the biological variations arising from predator-prey relationships, with cod's position far from mesozooplankton in the food chain, thus being less impacted; whereas NN forecasting incorporates various forms of uncertainty, including observation uncertainty, environmental shocks, and modeling errors. For the pelagic stock capelin, which is closer to primary production in the food web, the differences in uncertainty between the 2 models are smaller: the 2-SD bands by NNs are only slightly wider than those by Atlantis (see Fig. 6c,d).

The simulation outcomes from the 2 models show a high degree of similarity overall, indicating that NNs effectively capture the essential dynamics of the chosen system, especially the valuable cod stock, repre-

sented by Atlantis. While there are some discrepancies between the forecasted paths, these variations fall within an acceptable range for bioeconomic analysis. The general agreement between the 2 models strengthens the credibility of the NN representation, which is the central argument in our analyses.

#### 4. CONCLUDING REMARKS

In this study, we employed 3 analyses to progressively validate the effectiveness of our approach. The first analysis showcases that NNs as function representations can capture key dynamics of the subsystem and replicate the trajectories of the training data to a large extent. Stochasticity can be added to the dynamics in a straightforward and flexible manner. The second analysis illustrates that the trained NNs can efficiently evaluate various alternative policies, offering significantly reduced run time and greater ease in storage and application. The last analysis demonstrates that NNs are able to generalize to new sequences of harvest rates and produce almost identical trajectories as Atlantis for the commercially valuable cod stock. The finding that NNs can reasonably predict the behavior of the Atlantis model under a feedback rule is valuable and interesting because the underlying version of the Atlantis model cannot directly implement feedback harvest rules. Our findings suggest that the precision provided by NNs is adequate for high-level decision-making, and NNs can serve as relevant representations of large-scale ecosystem simulators.

The main challenge in applying NNs to represent ecosystem models is the need for abundant and relevant training data. In our analyses, we benefited from the effort put into the NoBa Atlantis model that provided us with synthetic data (Hansen et al. 2019b). In this regard, we emphasize that using synthetic data makes our networks vulnerable to weaknesses in the Atlantis model. Therefore, it is paramount that NNs are built upon data generated by simulator models that are comprehensive, cutting-edge, and state-of-the-art.

We promote the mathematical structure of feedforward NNs to represent ecosystem dynamics, potentially as representations of large-scale ecosystem simulators for various applications. Sustainable and ecosystem-based management decisions require an integrated and quantitative representation of complex and high-dimensional dynamic systems (Botsford et al. 1997, Levin et al. 2012). State-of-the-art ecosystem simulators, such as Atlantis-type models

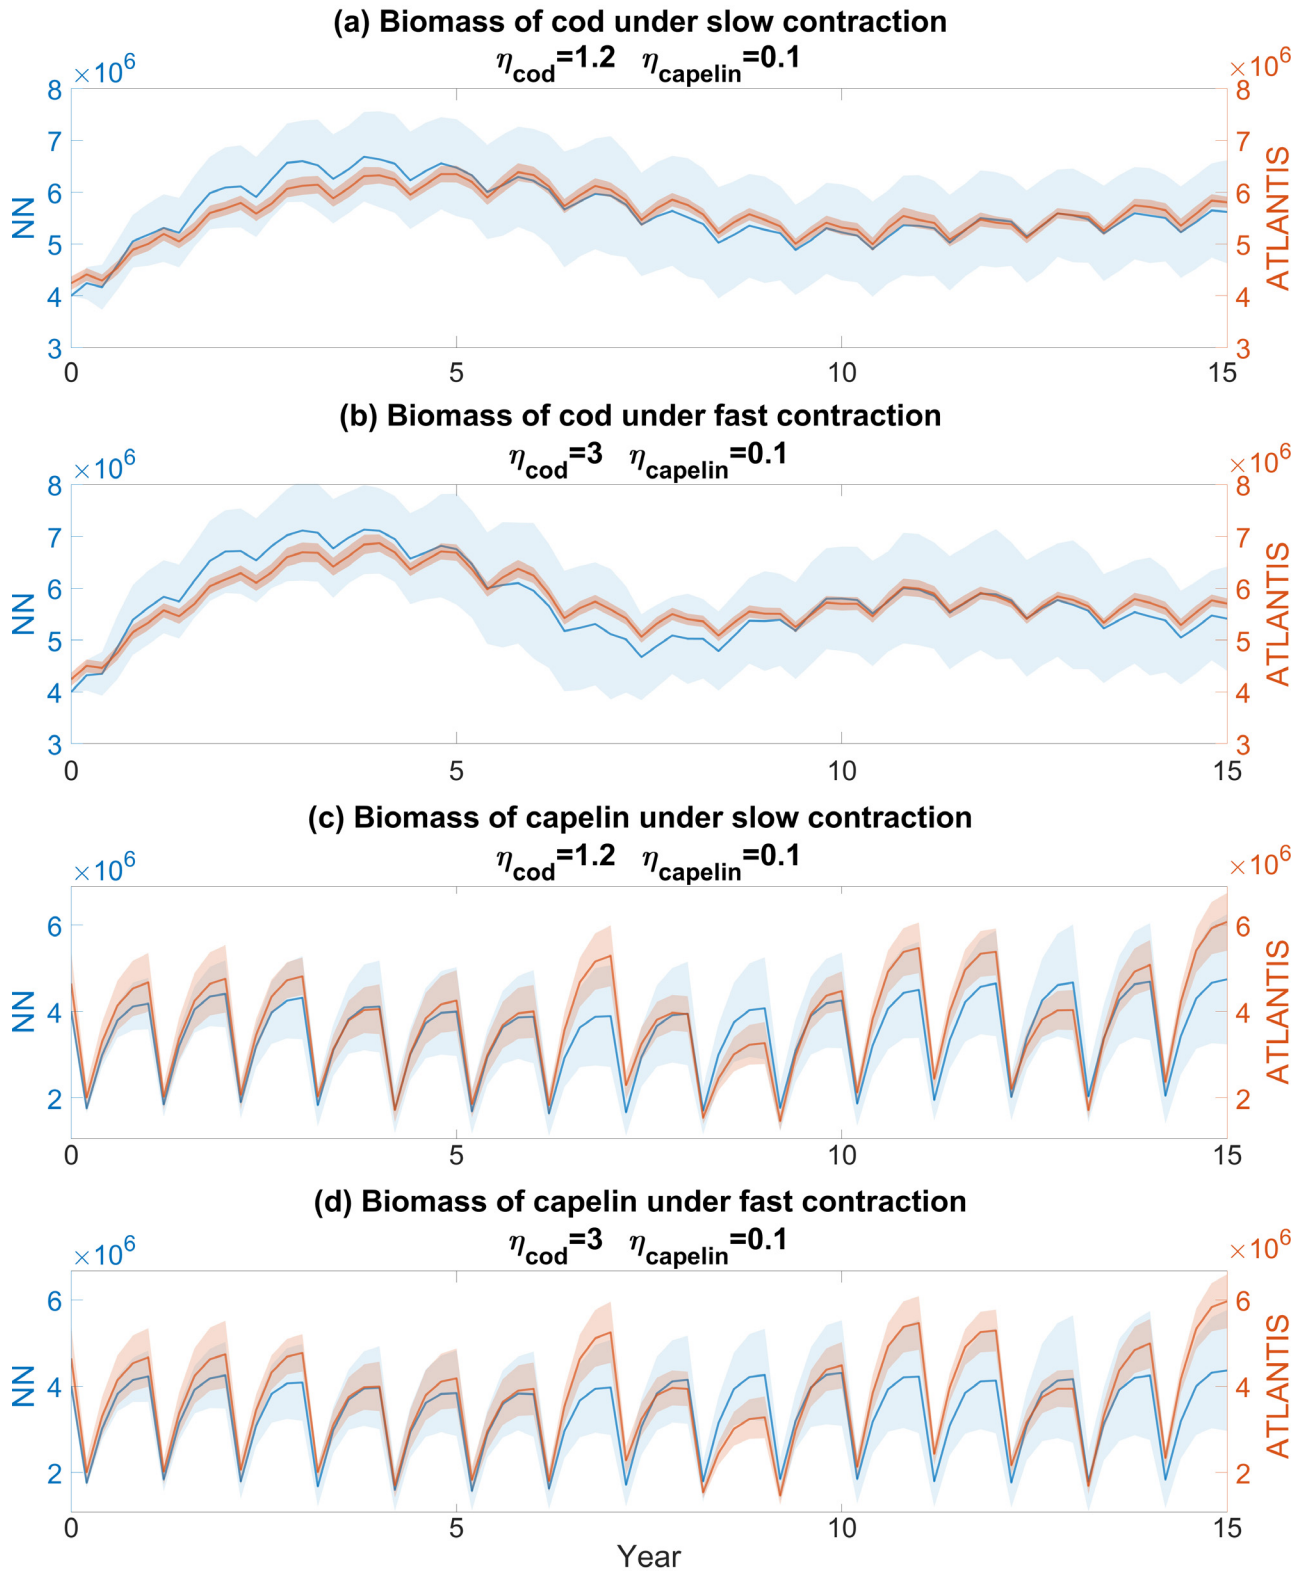


Fig. 6. Biomass level paths in tonnes of cod and capelin generated by stochastic NN forecasting (blue y-axis on the left) versus the Atlantis simulator (orange y-axis on the right) under the yearly updated adaptive policy with 2 contraction speeds  $\eta$ : (a,c) slow and (b,d) fast. Blue solid line: mean path of biomass generated by NNs. Orange solid line: mean path of biomass generated by Atlantis. Light blue bands: 2-SD of the 100 biomass paths generated by NNs. Light orange bands: 2-SD of the 14 biomass paths generated by Atlantis

(Fulton et al. 2011), are comprehensive representations of the knowledge and empirical experience regarding ecosystems, but require substantial computational resources (Hansen et al. 2019b) and are not particularly designed for dynamic decision analysis. NNs are universal nonlinear function approximators (Willcock et al. 2018) that satisfy the demands for efficiency and flexibility concerning the iterative algorithms used in bioeconomic analysis (Clark 1990, Sandal et al. 2021).

Earlier bioeconomic analyses predominantly relied on explicit closed-form models. The NN approach grants us freedom from the necessity to assume functional forms, which is an advantage given that our biological knowledge about ecosystems in general is still limited and developing. Further, by using synthetic data, we bypass problems with infrequent and incomplete ecosystem observations. To incorporate seasonal dynamics in the model often makes the bioeconomic problem more complex (Ni & Sandal 2019, Sandal et al. 2021), but it is important given the inherent seasonal variations in many natural systems (Kvamsdal et al. 2020a, Sandal et al. 2021). While it is not unachievable in closed-form models, using NNs to include seasonality is conceptually easier and does not require additional modeling work.

Uncertainty is another fundamental feature of ecosystem dynamics. Dealing with uncertainty involves several aspects. Two are: representing the uncertainty (modeling), and making decisions under uncertainty (evaluating solutions). The NNs can represent uncertainty in a similar fashion as stochastic differential equations and support decision analysis, thus being relevant to both aspects. However, with feedforward NNs, uncertainty cannot, to our knowledge, be included directly. In our setup, the networks represent the expected drift. Once the networks are trained, we can easily model uncertainties outside the network structures and specify the form and magnitude of the stochasticity. One idea is to use deviations between the network predictions and the training data to build an empirical distribution for the stochasticity. This should be addressed in future research.

The main contributions of our approach to modeling ecosystem dynamics are its extensive and potential applications, unlocking new doors and opportunities in this field of research. While many existing ecosystem models (Plagányi 2007) and some studies that apply NNs (Krasnopolsky et al. 2005, Russo et al. 2014, 2019) can perform various what-if analyses, the policies implemented in these models are usually limited by specified structures. Moving from what-if type

analyses to optimization enables thorough evaluation of a broader range of candidate policies. Even when we consider only what-if type analyses, it can be intractable for simulator models, which carry extensive details, to implement scenario projections defined on a higher level of aggregation. Our NN approach can aggregate to a desired aggregation level, initiate from any starting state, adopt policies in feedback form, and be coupled with iterative decision algorithms such as dynamic optimization. Bioeconomic analysis relies on transition equations such as Eq. (1), and if the trained NNs can represent the mapping function in the equation, our approach applies.

In conclusion, NNs serve as a bridge, facilitating more comprehensive decision-making with an expanded policy space while retaining the empirical knowledge embedded in ecosystem simulators such as NoBa Atlantis. Thus, our application of NNs builds on and leverages the extensive development of ecosystem models to progress ecosystem-based management (Link 2010, Link et al. 2020) and ecosystem wealth assessments (Yun et al. 2017, Kvamsdal et al. 2020b) that underpin a sustainable development of the ocean economy (Fenichel et al. 2020). Future work along this line of research also includes applications of our NN approach to decision science frameworks, such as dynamic programming and reinforcement learning for natural resource management (Malo et al. 2021, Tahvonen et al. 2022).

*Acknowledgments.* This work was supported by the Research Council of Norway (Norges Forskningsråd) (grant numbers 302197 and 325665).

#### LITERATURE CITED

- Aanestad S, Sandal LK, Eide A (2007) Optimal fishing policy for two species in a three-species predator–prey model: the case of capelin, cod and juvenile herring in the Barents Sea. Working paper, Norwegian School of Economics and Business Administration, Department of Finance and Management Science
- ✦ Audzijonyte A, Pethybridge H, Porobic J, Gorton R, Kaplan I, Fulton EA (2019) Atlantis: a spatially explicit end-to-end marine ecosystem model with dynamically integrated physics, ecology and socio-economic modules. *Methods Ecol Evol* 10:1814–1819
- ✦ Botsford LW, Castilla JC, Peterson CH (1997) The management of fisheries and marine ecosystems. *Science* 277: 509–515
- ✦ Burkett VR, Wilcox DA, Stottlemeyer R, Barrow W and others (2005) Nonlinear dynamics in ecosystem response to climatic change: case studies and policy implications. *Ecol Complex* 2:357–394
- ✦ Cai Y, Lontzek TS (2019) The social cost of carbon with economic and climate risks. *J Polit Econ* 127:2684–2734

- Clark CW (1990) *Mathematical bioeconomics: the optimal management of renewable resources*, 2nd Edn. Wiley-Interscience, New York
- Crépin AS, Norberg J, Mäler KG (2011) Coupled economic-ecological systems with slow and fast dynamics—modelling and analysis method. *Ecol Econ* 70:1448–1458.
- Deneu B, Joly A, Bonnet P, Servajean M, Munoz F (2022) Very high resolution species distribution modeling based on remote sensing imagery: how to capture fine-grained and large-scale vegetation ecology with convolutional neural networks? *Front Plant Sci* 13:839279
- Durant JM, Hjermann DO, Sabarros PS, Stenseth NC (2008) Northeast Arctic cod population persistence in the Lofoten-Barents Sea system under fishing. *Ecol Appl* 18:662–669
- Ekerhovd NA, Kvamsdal SF (2017) Up the ante on bioeconomic submodels of marine food webs: a data assimilation-based approach. *Ecol Econ* 131:250–261
- Fenichel EP, Addicott ET, Grimsrud KM, Lange GM, Porras I, Milligan B (2020) Modifying national accounts for sustainable ocean development. *Nat Sustain* 3:889–895
- Fulton EA, Link JS, Kaplan IC, Savina-Rolland M and others (2011) Lessons in modelling and management of marine ecosystems: the Atlantis experience. *Fish Fish* 12:171–188
- Goodfellow I, Bengio Y, Courville A (2016) *Deep learning*. The MIT Press, Cambridge, MA
- Hansen C, Drinkwater KF, Jähkel A, Fulton EA, Gorton R, Skern-Mauritzen M (2019a) Sensitivity of the Norwegian and Barents Sea Atlantis end-to-end ecosystem model to parameter perturbations of key species. *PLOS ONE* 14:e0210419
- Hansen C, Nash RDM, Drinkwater KF, Hjøllø SS (2019b) Management scenarios under climate change—a study of the Nordic and Barents Seas. *Front Mar Sci* 6:668
- Helstad K (2001) Management of herring, capelin and cod in the greater Barents Sea—economic optimal management from a Norwegian point of view. *Proc International Institute of Fisheries Economics & Trade (IIFET)*, July 10–14, 2000. IIFET, Corvallis, OR
- Holt J, Icarus Allen J, Anderson TR, Brewin R and others (2014) Challenges in integrative approaches to modelling the marine ecosystems of the North Atlantic: physics to fish and coasts to ocean. *Prog Oceanogr* 129:285–313
- Hornik K, Stinchcombe M, White H (1989) Multilayer feedforward networks are universal approximators. *Neural Netw* 2:359–366
- Huang L, Smith MD (2014) The dynamic efficiency costs of common-pool resource exploitation. *Am Econ Rev* 104:4071–4103
- Hutter F, Lücke J, Schmidt-Thieme L (2015) Beyond manual tuning of hyperparameters. *Künstl Intell* 29:329–337
- ICES (2018) *Publication Reports—Introduction to Advice 2018*. ICES, Copenhagen
- ICES (2020) *Arctic Fisheries Working Group (AFWG) (report)*. ICES Scientific Reports. ICES, Copenhagen
- ICES (2021) *Cod (*Gadus morhua*) in subareas 1 and 2 (Northeast Arctic)*. ICES Advice: Recurrent Advice. ICES, Copenhagen
- Javid AM, Das S, Skoglund M, Chatterjee S (2020) A ReLU dense layer to improve the performance of neural networks. In: *ICASSP 2021–2021, IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, p 2810–2814
- Joseph MB (2020) Neural hierarchical models of ecological populations. *Ecol Lett* 23:734–747
- Krasnopolsky VM, Fox-Rabinovitz MS, Chalikov DV (2005) New approach to calculation of atmospheric model physics: accurate and fast neural network emulation of longwave radiation in a climate model. *Mon Weather Rev* 133:1370–1383
- Kvamsdal SF, Sandal LK (2015) The ensemble Kalman filter for multidimensional bioeconomic models. *Nat Resour Model* 28:321–347.
- Kvamsdal SF, Eide A, Ekerhovd NA, Enberg K and others (2016) Harvest control rules in modern fisheries management. *Elem Sci Anthr* 4:000114
- Kvamsdal S, Maroto JM, Morán M, Sandal LK (2020a) Bioeconomic modeling of seasonal fisheries. *Eur J Oper Res* 281:332–340
- Kvamsdal SF, Sandal LK, Poudel D (2020b) Ecosystem wealth in the Barents Sea. *Ecol Econ* 171:106602
- Lek S, Guégan JF (1999) Artificial neural networks as a tool in ecological modelling: an introduction. *Ecol Modell* 120:65–73
- Levin K, Cashore B, Bernstein S, Auld G (2012) Overcoming the tragedy of super wicked problems: constraining our future selves to ameliorate global climate change. *Policy Sci* 45:123–152
- Levin S, Xepapadeas T, Crépin AS, Norberg J and others (2013) Social-ecological systems as complex adaptive systems: modeling and policy implications. *Environ Dev Econ* 18:111–132
- Link JS (2010) *Ecosystem-based fisheries management: confronting tradeoffs*. Cambridge University Press, Cambridge
- Link JS, Huse G, Gaichas S, Marshak AR (2020) Changing how we approach fisheries: a first attempt at an operational framework for ecosystem approaches to fisheries management. *Fish Fish* 21:393–434
- Lu F, Chen Z, Liu W, Shao H (2016) Modeling chlorophyll-a concentrations using an artificial neural network for precisely eco-restoring lake basin. *Ecol Eng* 95:422–429
- Malde K, Handegard NO, Eikvil L, Salberg AB (2020) Machine intelligence and the data-driven future of marine science. *ICES J Mar Sci* 77:1274–1285
- Malo P, Tahvonen O, Suominen A, Back P, Viitasaari L (2021) Reinforcement learning in optimizing forest management. *Can J For Res* 51:1393–1409
- May RM, Beddington JR, Clark CW, Holt SJ, Laws RM (1979) Management of multispecies fisheries. *Science* 205:267–277
- Møller MF (1993) A scaled conjugate gradient algorithm for fast supervised learning. *Neural Netw* 6:525–533
- Ni Y, Sandal LK (2019) Seasonality matters: a multi-season, multi-state dynamic optimization in fisheries. *Eur J Oper Res* 275:648–658
- Olsen E, Fay G, Gaichas S, Gamble R, Lucey S, Link JS (2016) Ecosystem model skill assessment. *Yes we can!* *PLOS ONE* 11:e0146467
- Plagányi ÉE (2007) *Models for an ecosystem approach to fisheries*. Food and Agricultural Organization, Rome
- Planque B (2016) Projecting the future state of marine ecosystems, ‘la grande illusion’? *ICES J Mar Sci* 73:204–208
- Poudel D, Sandal LK, Kvamsdal SF (2015) Stochastically induced critical depensation and risk of stock collapse. *Mar Resour Econ* 30:297–313
- Rammer W, Seidl R (2019) Harnessing deep learning in ecology: an example predicting bark beetle outbreaks. *Front Plant Sci* 10

- ✦ Russo T, Parisi A, Garofalo G, Gristina M, Cataudella S, Fiorentino F (2014) SMART: a spatially explicit bio-economic model for assessing and managing demersal fisheries, with an application to Italian trawlers in the Strait of Sicily. *PLOS ONE* 9:e86222
- ✦ Russo T, D'Andrea L, Franceschini S, Accadia P and others (2019) Simulating the effects of alternative management measures of trawl fisheries in the central Mediterranean Sea: application of a multi-species bio-economic modeling approach. *Front Mar Sci* 6:542
- ✦ Sandal LK, Steinshamn SI (1997) A stochastic feedback model for optimal management of renewable resources. *Nat Resour Model* 10:31–52
- ✦ Sandal LK, Kvamsdal SF, Maroto JM, Morán M (2021) A contraction approach to dynamic optimization problems. *PLOS ONE* 16:e0260257
- ✦ Sandø AB, Melsom A, Budgell WP (2014) Downscaling IPCC control run and future scenario with focus on the Barents Sea. *Ocean Dyn* 64:927–949
- ✦ Scowen M, Athanasiadis IN, Bullock JM, Eigenbrod F, Willcock S (2021) The current and future uses of machine learning in ecosystem service research. *Sci Total Environ* 799:149263
- ✦ Shchepetkin AF, McWilliams JC (2005) The regional oceanic modeling system (ROMS): a split-explicit, free-surface, topography-following-coordinate oceanic model. *Ocean Model* 9:347–404
- ✦ Skogen MD, Hjøllø SS, Sandø AB, Tjiputra J (2018) Future ecosystem changes in the Northeast Atlantic: a comparison between a global and a regional model system. *ICES J Mar Sci* 75:2355–2369
- ✦ Smith MD (2012) The new fisheries economics: incentives across many margins. *Annu Rev Resour Econ* 4:379–402.
- ✦ Springborn MR, Faig A (2019) Moving forward: a simulation-based approach for solving dynamic resource management problems. *Mar Resour Econ* 34:199–224
- ✦ Tahvonen O, Suominen A, Malo P, Viitasaari L, Parkatti VP (2022) Optimizing high-dimensional stochastic forestry via reinforcement learning. *J Econ Dyn Control* 145:104553
- ✦ Ussif AAM, Sumaila UR (2005) Modeling the dynamics of regulated resource systems: a fishery example. *Ecol Econ* 52:469–479
- ✦ Willcock S, Martínez-López J, Hooftman DAP, Bagstad KJ and others (2018) Machine learning for ecosystem services. *Ecosyst Serv* 33:165–174
- ✦ Yun SD, Hutniczak B, Abbott JK, Fenichel EP (2017) Ecosystem-based management and the wealth of ecosystems. *Proc Natl Acad Sci USA* 114:6539–6544

*Editorial responsibility: Ingrid van Putten,  
Hobart, Tasmania, Australia*  
*Reviewed by: T. Russo and 1 anonymous referee*

*Submitted: October 26, 2022*  
*Accepted: June 19, 2023*  
*Proofs received from author(s): August 4, 2023*