UCLA

Computer Science

Bio-Inspired Simulation With Learning-Based Automatic Motion Control

Xiao (Steven) Zeng

Dec 6, 2022

Bio Snippet

- 4th year Ph.D. candidate
- Advisor: Demetri Terzopoulos
- Research:
  - Muscle-driven biomechanical human animation and neuromuscular control
  - Adaptive simulation of fish schooling behavior
- Relevant publications:
  - Zeng, X. S., Ishiwaka, Y., Ogawa, S., Westwater, D. M., Tone, T., and Nakada, M. (2022). DeepFoids: Adaptive Bio-Inspired Fish Simulation with Deep Reinforcement Learning. *Advances in Neural Information Processing System (NeurIPS) 2022*
  - Zeng, X. S., Dwarakanath, S., Lu, W., Nakada, M., and Terzopoulos, D. (2021). Neuromuscular Control of the Face-Head-Neck Biomechanical Complex with Learning-Based Expression Transfer from Images and Videos. In *The 16th International Symposium on Visual Computing (ISVC 2021)*
  - Zeng, X. S., Dwarakanath, S., Lu, W., Nakada, M., and Terzopoulos, D. (2021). Facial Expression Transfer from Video Via Deep Learning. In *The 20th ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA 2021)*
  - Ishiwaka, Y., Zeng, X. S., Eastman, M. L., Kakazu, S., Gross, S., Mizutani, R., and Nakada, M. (2021). Foids: bio-inspired fish simulation for generating synthetic datasets. *ACM Transactions on Graphics (TOG)*

Outline

I. Motivation and Objectives
II. Related Work
    A. Muscle-driven Facial Animation
    B. Fish Schooling Simulation
III. Face-Head-Neck Biomechanical Model
    A. Overview
    B. Musculoskeletal Model
    C. Control
    D. Result
IV. Adaptive Bio-inspired Fish Simulation
    A. Overview
    B. Fish Simulation
    C. Control
    D. Result
V. Summery and Future Work

Motivation

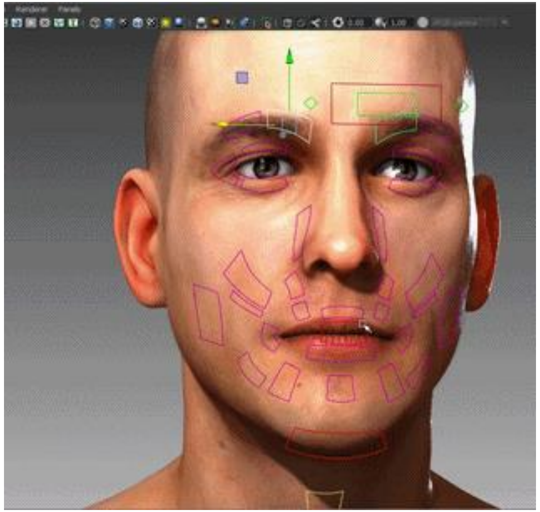

FACS-based
Facial Rigging

Fish Schooling
Animation

Objectives

Generate animation in a fully autonomous manner at a micro (muscle control) and macro (flocking control) levels of abstraction

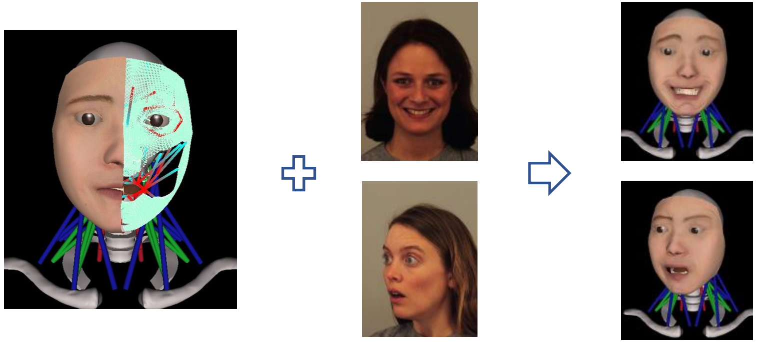

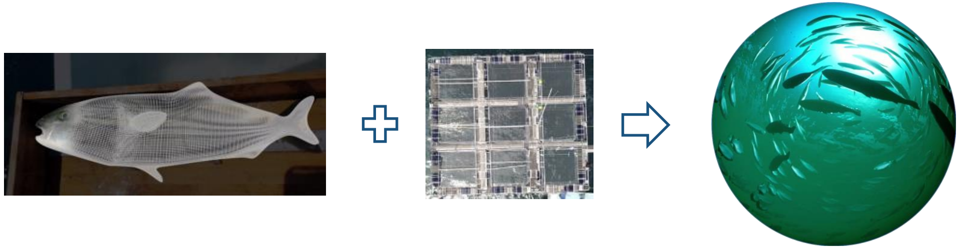

Related Work - Muscle-driven Facial Animation

Biomechanical face

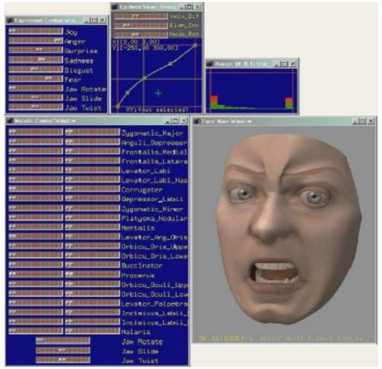

- [Terzopoulos and Lee 2004]
- [Lee et al. 1995]
- [Wu et al. 2014]

...

Musculoskeletal head-neck system

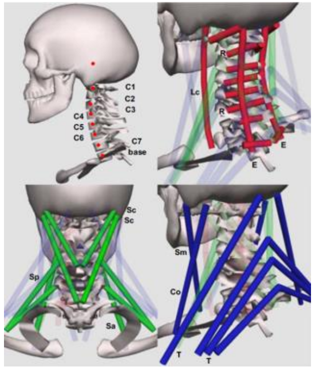

- [Lee and Terzopoulos 2006]

...

Expression transfer

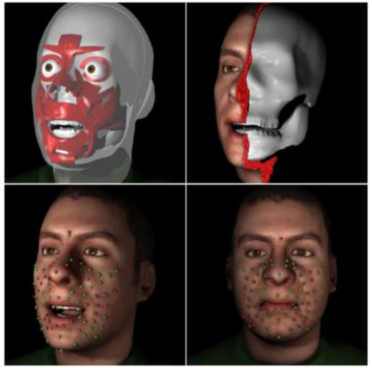

- [Sifakis et al. 2005]
- [Weise et al. 2011]
- [Zhang et al. 2020]

...

FACS

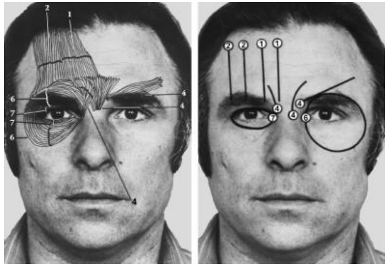

- [Ekman et al. 2002]
- [Cohn et al. 2007]

...

Related Work – Fish Schooling Simulation

Fish Simulation

- [Reynolds, 1987]
- [Tu and Terzopoulos 1994]
- [Satoi et al. 2016]
- [Ishiwaka et al. 2021]

...

DRL-Driven Swimming Simulation

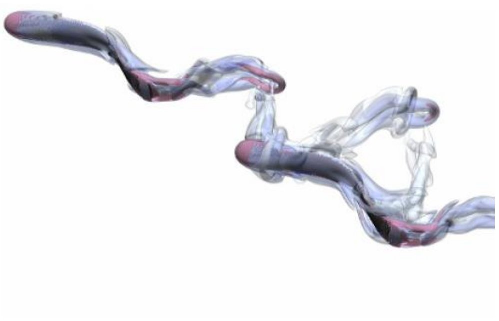

- [Verma et al. 2018]
- [Zhu et al. 2021]
- [Sunehag et al. 2019]
- [Lee et al. 2018]

...

Biological Background

- [Reynolds and Casterlin 1979]
- [Jensen et al. 1989]
- [Juell 1995]
- [Oppedal et al. 2011]
- [Miller et al. 2017]
- [Sakakura and Tsukamoto 1998]

...

Face-Head-Neck Biomechanical Model

Overview

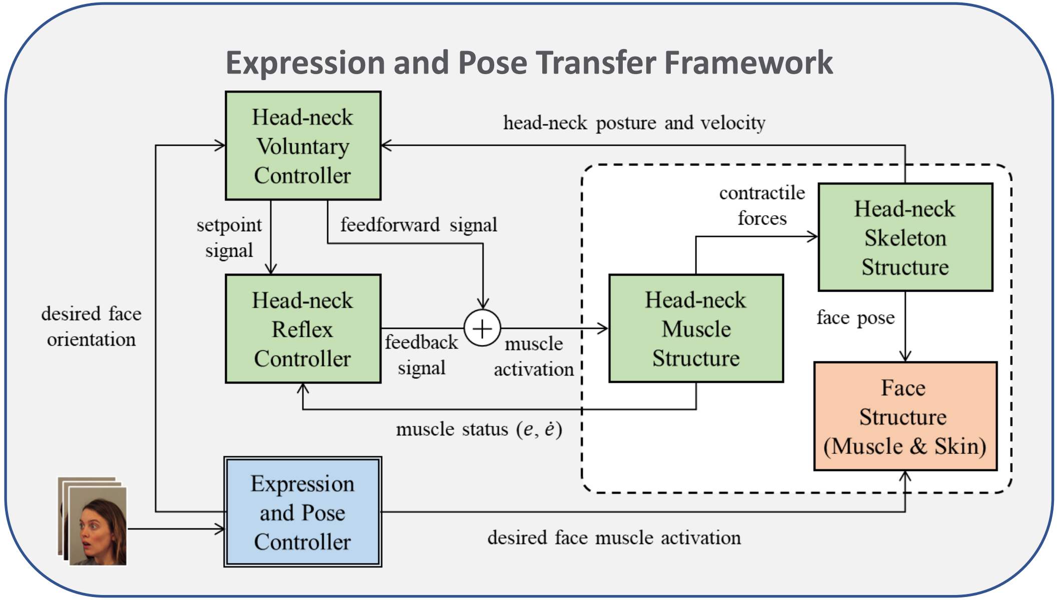

Our framework has three components:

- Biomechanical musculoskeletal head-neck system (green)
- Biomechanical face simulator (red)
- Facial expression and head pose controller (blue)

Musculoskeletal Model - Head-neck System

- Anatomically consistent
- Voluntary controller
  - Outputs muscle activation & setpoint signals to achieve desired head orientations
- Reflex controller
  - Adjusts muscle activation according to the setpoint signal, muscle strain & strain rate
  - Passes the activation to muscles
- Muscles receives activation and output contractile force to pull the skeleton
- Has 3 rotational degrees of freedom
  - Pitch, roll & yaw

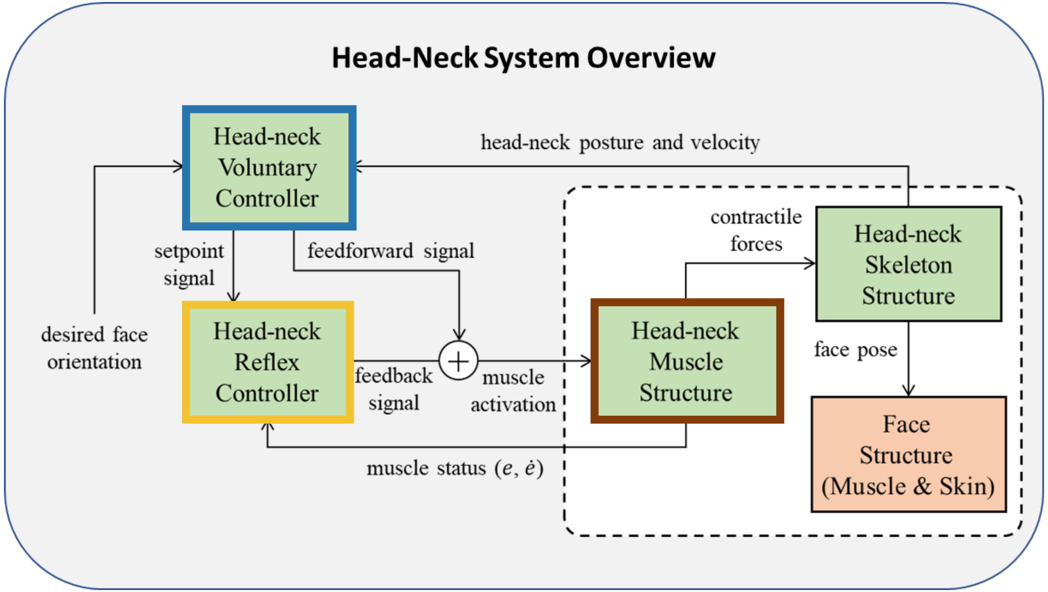

UCLA

Musculoskeletal Model - Head-neck System

- Neck Muscles are attached to bones via connection points, and are modeled using a modified version of the Hill-type actuator
- Muscle force computation:

$$f_m = f_P + f_C$$

f_P: a passive restoring force due to muscle elasticity

f_C: a proactive contractile force, proportional to the activation level between [0, 1] and influenced by the force-length relation and force-velocity relation of the muscle model

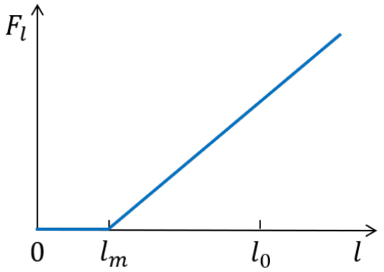

(a) Force-length relationship

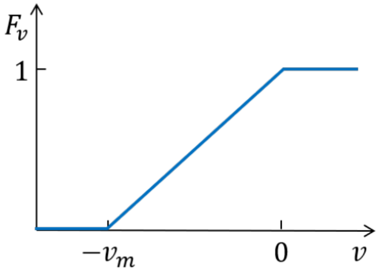

(b) Force-velocity relationship

Musculoskeletal Model - Biomechanical Face

- Improved upon the work of Lee and Terzopoulos (2006)
- A physics-based, real-time simulation system
- Contains skull, muscle, and tissue layers
- Tissue is a lattice of uniaxial viscoelastic units, arranged in a multilayered prismatic structure containing fascia, dermal-fatty, and epidermal layers
- Synthesizes 6 basic expressions and their combinations
  - Joy
  - Anger
  - Surprise
  - Sadness
  - Disgust
  - Fear

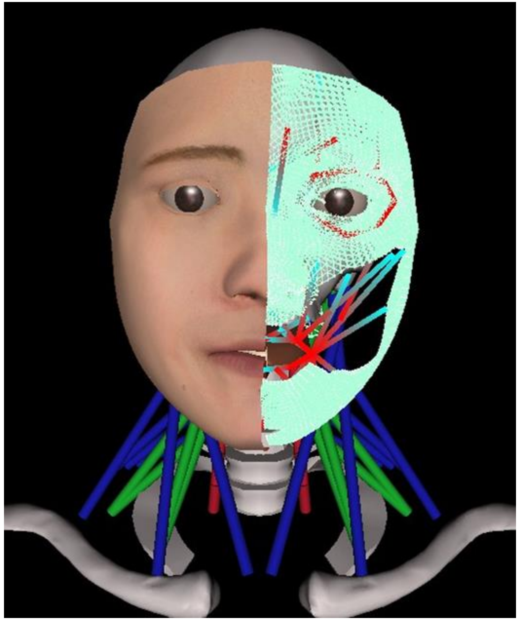

Demo of the face-head-neck system

Head-Neck System

Control - Expression and Head Pose Transfer

- The control pipeline consists of
  - OpenFace
  - Offline Learning Module
  - Online Transfer Module
- Trained offline using synthetic dataset generated by the biomechanical face model itself
- The expression and head pose transfer controller takes an image as input and produces muscle activations and jaw rotation as output

OpenFace

- OpenFace (Baltrusaitis et al., 2018) is a toolkit which, given an image of the face, is capable of:
  - FACS based Action Unit Estimation (presence and intensity)
  - Facial Landmark Tracking
  - Eye Gaze Estimation
  - Head Orientation Estimation
- OpenFace is built using DNNs and detects AUs using linear kernel Support Vector Machines

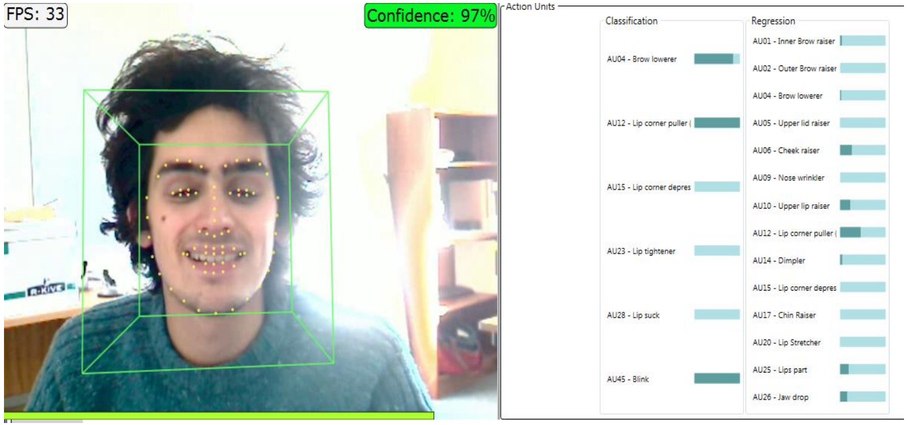

Offline Learning Module

- The offline learning module uses the data generated from the biomechanical model itself
- Training data generated by random expressions created by some variation of muscle activations
  - A basic scheme for the muscle weight is decided using visual analysis (only once)
- The network is trained on training pairs which consist of Action Units as input and muscle activations as output

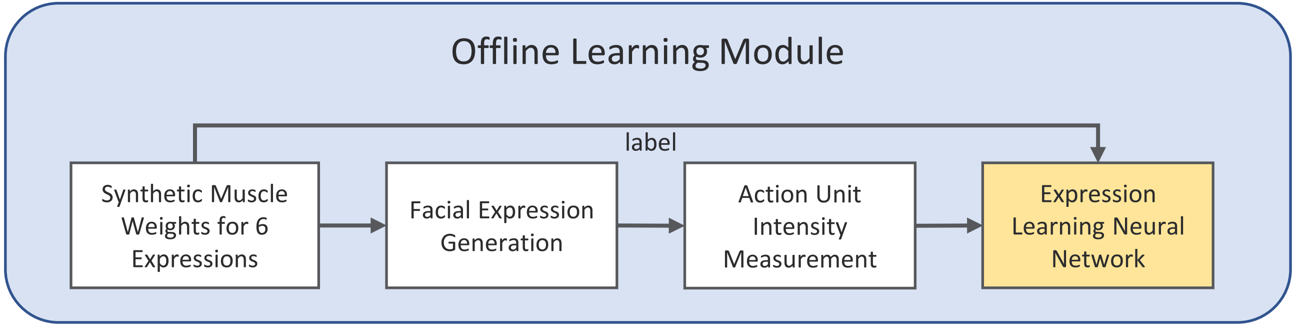

Online Transfer Module

- Once trained, the neural network can be readily used for online expression transfer
- The neural network transfers the AU values received by OpenFace into muscle activations

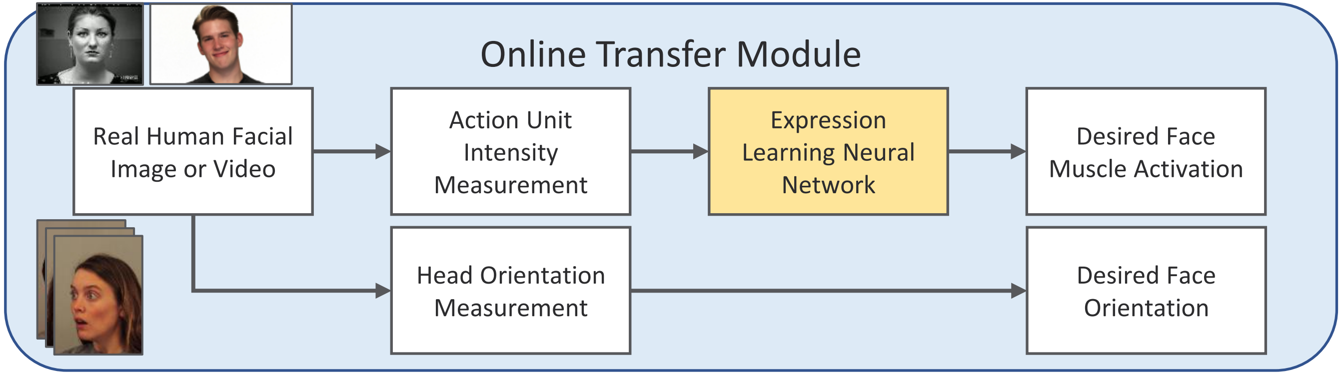

Neural Network Architecture

- A neural network with 4 hidden layers with 100 nodes each
- Architecture was decide after comparing the MSE of networks with different number of layers
- The input dimensions are equal to the number of AUs output by OpenFace
- The output dimensions are the number of muscles and jaw activations

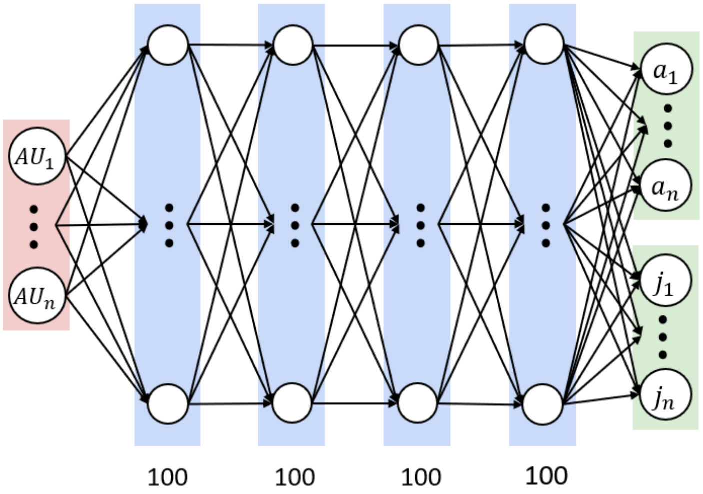

Transfer Result (KDEF Dataset)

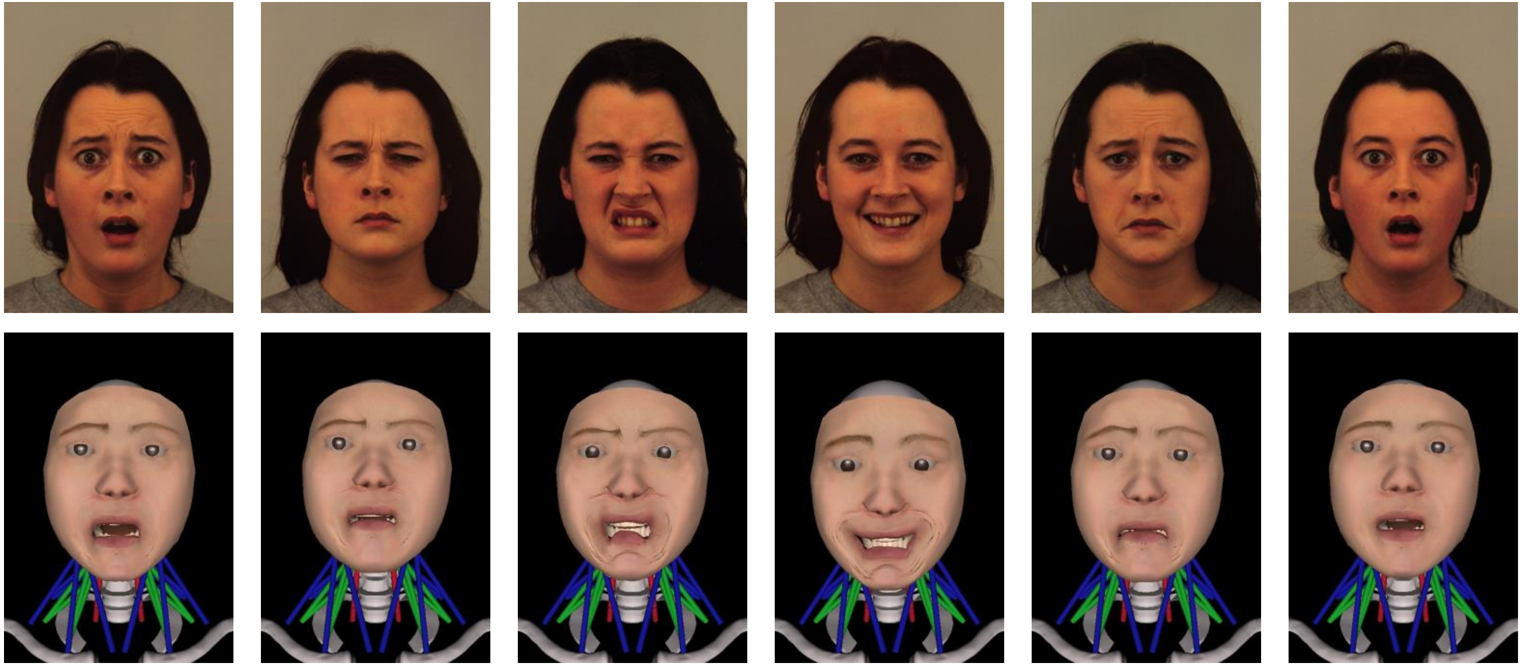

Fear Anger Disgust Joy Sadness Surprise

Transfer Result with Orientation (KDEF Dataset)

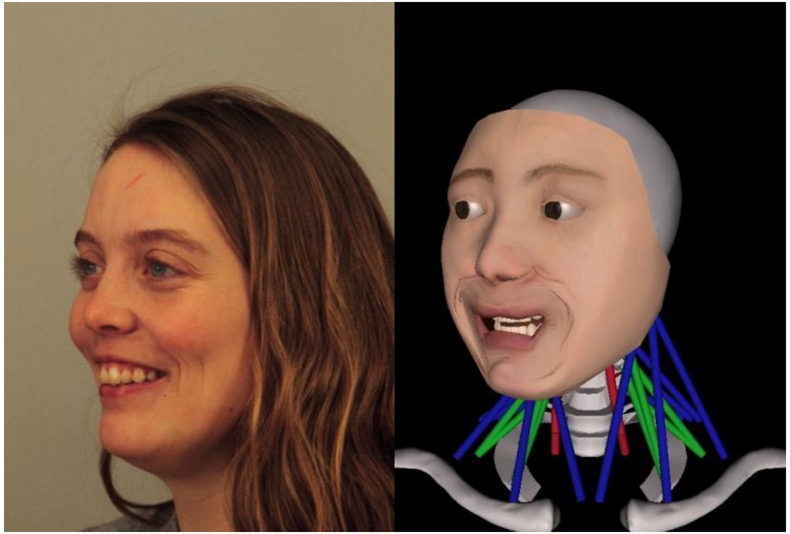

Joy + Right Turn

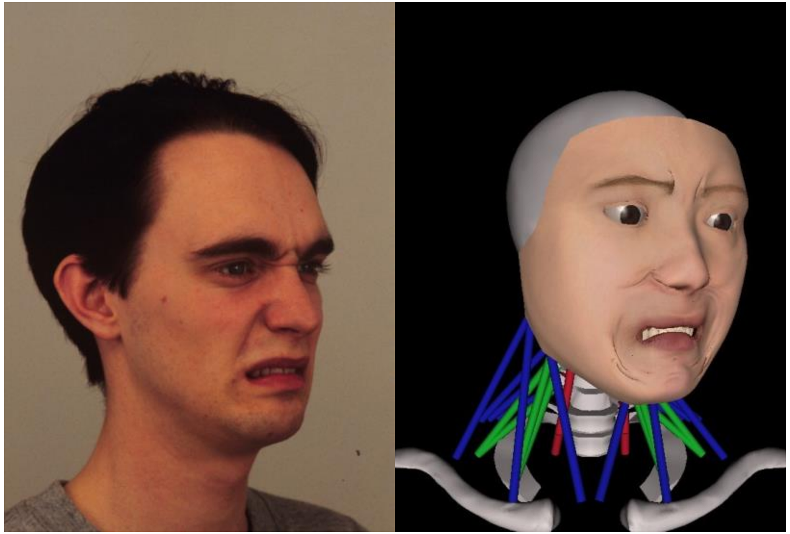

Disgust + Left Turn

Demo of video transfer result

Video Transfer Result

Adaptive Bio-inspired Fish Simulation

Overview

A simulated ecosystem of schooling fish, named DeepFoids, utilizing computer graphics and deep reinforcement learning for reproducing realistic and various schooling behaviors autonomously

- Improved upon the Foids model (Ishiwaka et al., 2021)
- Generates biologically valid synthetic data to train an automated fish counting system

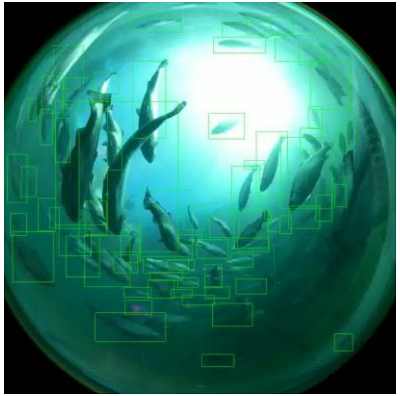

Biological Features

$$\Delta v_{temp} = \begin{cases} \dfrac{T_{lpref} - T}{T_{lpref} - zT_{lsteep}} [0, 1, 0] & if\ T \le T_{lpref} \\ \dfrac{T - T_{hpref}}{T_{hsteep} - T_{hpref}} [0, -1, 0] & if\ T \ge T_{hpref} \end{cases}$$

$$\Delta v_{light} = \begin{cases} \dfrac{I_{lpref} - I}{I_{lpref} - I_{lsteep}} [0, 1, 0] & if\ I \le I_{lpref} \\ \dfrac{I - I_{hpref}}{I_{hsteep} - I_{hpref}} [0, -1, 0] & if\ I \ge I_{hpref} \end{cases}$$

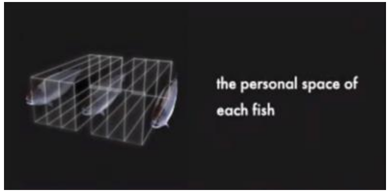

Control

Deep Reinforcement Learning (DRL)-Driven Behavior Model

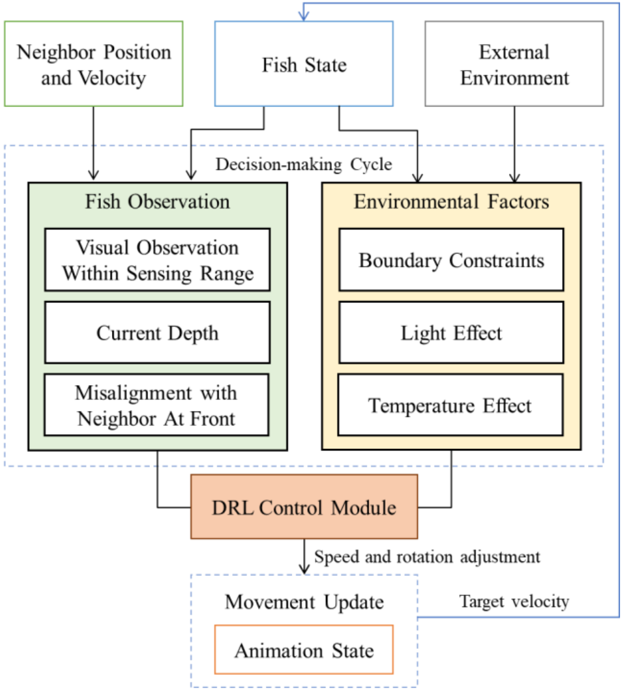

DRL-Driven Behavior Model

- Observation

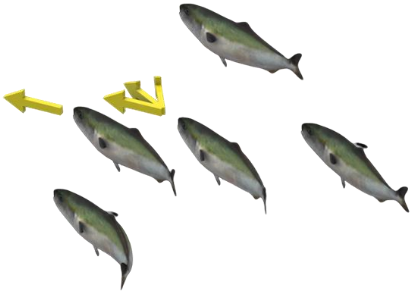

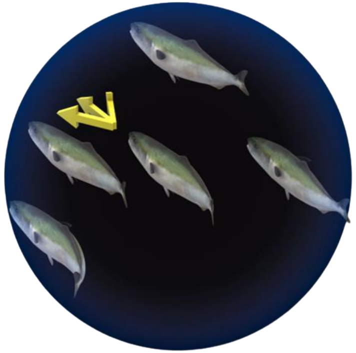

- Action

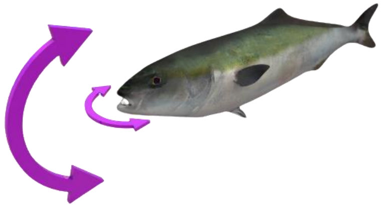

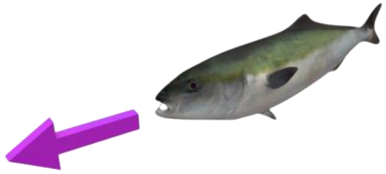

DRL-Driven Behavior Model

Reward:

$$r_t = r_t^{BC} + r_t^{NC} + r_t^{BD} + r_t^{ND} + r_t^{E} + r_t^{M} + r_t^{C}$$

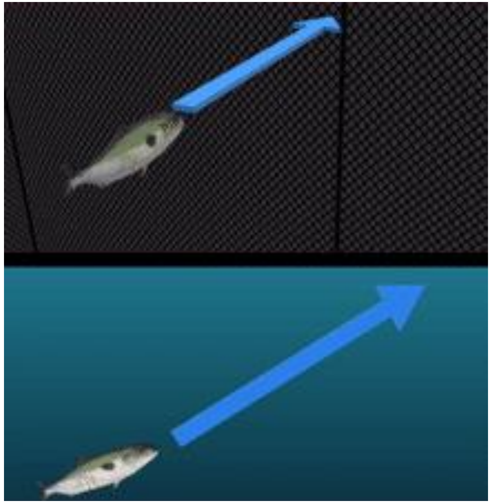

r_t^{BC} - A penalty given when a fish collides with the cage or water surface

r_t^{NC} - A penalty given when a fish collides with neighbor fish

r_t^{BD} - A reward given when a fish avoids the cage and water surface

r_t^{ND} - A reward given when a fish stays close to neighbor fish

r_t^{E} - A penalty given when a fish expends energy

r_t^{M} - A reward given when a fish swim faster than a minimum speed w/o abrupt change in depth, or a penalty given when a fish abruptly changes its depth

r_t^{C} - A reward given when a dominant fish attacks its target, or a penalty given when a subordinate fish is being attacked

DRL-Driven Behavior Model

Reward:

$$r_t = r_t^{BC} + r_t^{NC} + r_t^{BD} + r_t^{ND} + r_t^{E} + r_t^{M} + r_t^{C}$$

r_t^{BC} - A penalty given when a fish collides with the cage or water surface

r_t^{NC} - A penalty given when a fish collides with neighbor fish

r_t^{BD} - A reward given when a fish avoids the cage and water surface

r_t^{ND} - A reward given when a fish stays close to neighbor fish

r_t^{E} - A penalty given when a fish expends energy

r_t^{M} - A reward given when a fish swim faster than a minimum speed w/o abrupt change in depth,
or a penalty given when a fish abruptly changes its depth

r_t^{C} - A reward given when a dominant fish attacks its target,
or a penalty given when a subordinate fish is being attacked

UCLA

DRL-Driven Behavior Model

Reward:

$$r_t = r_t^{BC} + r_t^{NC} + r_t^{BD} + r_t^{ND} + r_t^{E} + r_t^{M} + r_t^{C}$$

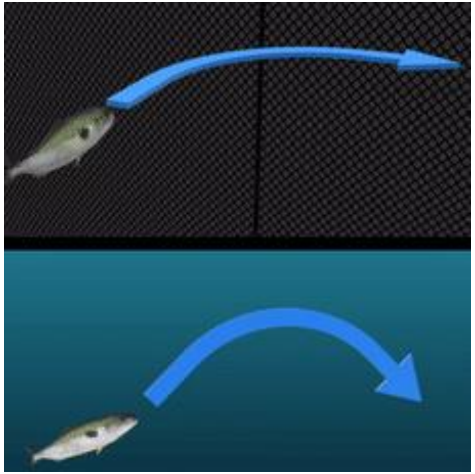

r_t^{BC} - A penalty given when a fish collides with the cage or water surface

r_t^{NC} - A penalty given when a fish collides with neighbor fish

r_t^{BD} - A reward given when a fish avoids the cage and water surface

r_t^{ND} - A reward given when a fish stays close to neighbor fish

r_t^{E} - A penalty given when a fish expends energy

r_t^{M} - A reward given when a fish swim faster than a minimum speed w/o abrupt change in depth, or a penalty given when a fish abruptly changes its depth

r_t^{C} - A reward given when a dominant fish attacks its target, or a penalty given when a subordinate fish is being attacked

DRL-Driven Behavior Model

Reward:

$$r_t = r_t^{BC} + r_t^{NC} + r_t^{BD} + r_t^{ND} + r_t^{E} + r_t^{M} + r_t^{C}$$

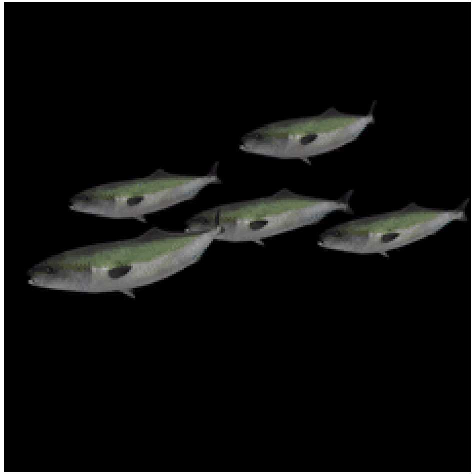

r_t^{BC} - A penalty given when a fish collides with the cage or water surface

r_t^{NC} - A penalty given when a fish collides with neighbor fish

r_t^{BD} - A reward given when a fish avoids the cage and water surface

r_t^{ND} - A reward given when a fish stays close to neighbor fish

r_t^{E} - A penalty given when a fish expends energy

r_t^{M} - A reward given when a fish swim faster than a minimum speed w/o abrupt change in depth,

or a penalty given when a fish abruptly changes its depth

r_t^{C} - A reward given when a dominant fish attacks its target,

or a penalty given when a subordinate fish is being attacked

DRL-Driven Behavior Model

Reward:

$$r_t = r_t^{BC} + r_t^{NC} + r_t^{BD} + r_t^{ND} + r_t^{E} + r_t^{M} + r_t^{C}$$

r_t^{BC} - A penalty given when a fish collides with the cage or water surface

r_t^{NC} - A penalty given when a fish collides with neighbor fish

r_t^{BD} - A reward given when a fish avoids the cage and water surface

r_t^{ND} - A reward given when a fish stays close to neighbor fish

r_t^{E} - A penalty given when a fish expends energy

r_t^{M} - A reward given when a fish swim faster than a minimum speed w/o abrupt change in depth,

or a penalty given when a fish abruptly changes its depth

r_t^{C} - A reward given when a dominant fish attacks its target,

or a penalty given when a subordinate fish is being attacked

UCLA

DRL-Driven Behavior Model

Reward:

$$r_t = r_t^{BC} + r_t^{NC} + r_t^{BD} + r_t^{ND} + r_t^{E} + r_t^{M} + r_t^{C}$$

r_t^{BC} - A penalty given when a fish collides with the cage or water surface

r_t^{NC} - A penalty given when a fish collides with neighbor fish

r_t^{BD} - A reward given when a fish avoids the cage and water surface

r_t^{ND} - A reward given when a fish stays close to neighbor fish

r_t^{E} - A penalty given when a fish expends energy

r_t^{M} - A reward given when a fish swim faster than a minimum speed w/o abrupt change in depth,

or a penalty given when a fish abruptly changes its depth

r_t^{C} - A reward given when a dominant fish attacks its target,

or a penalty given when a subordinate fish is being attacked

DRL-Driven Behavior Model

Reward:

$$r_t = r_t^{BC} + r_t^{NC} + r_t^{BD} + r_t^{ND} + r_t^{E} + r_t^{M} + r_t^{C}$$

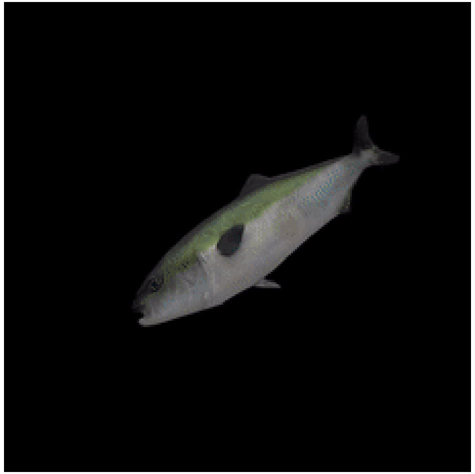

r_t^{BC} - A penalty given when a fish collides with the cage or water surface

r_t^{NC} - A penalty given when a fish collides with neighbor fish

r_t^{BD} - A reward given when a fish avoids the cage and water surface

r_t^{ND} - A reward given when a fish stays close to neighbor fish

r_t^{E} - A penalty given when a fish expends energy

r_t^{M} - A reward given when a fish swim faster than a minimum speed w/o abrupt change in depth, or a penalty given when a fish abruptly changes its depth

r_t^{C} - A reward given when a dominant fish attacks its target, or a penalty given when a subordinate fish is being attacked

Experiment Setting

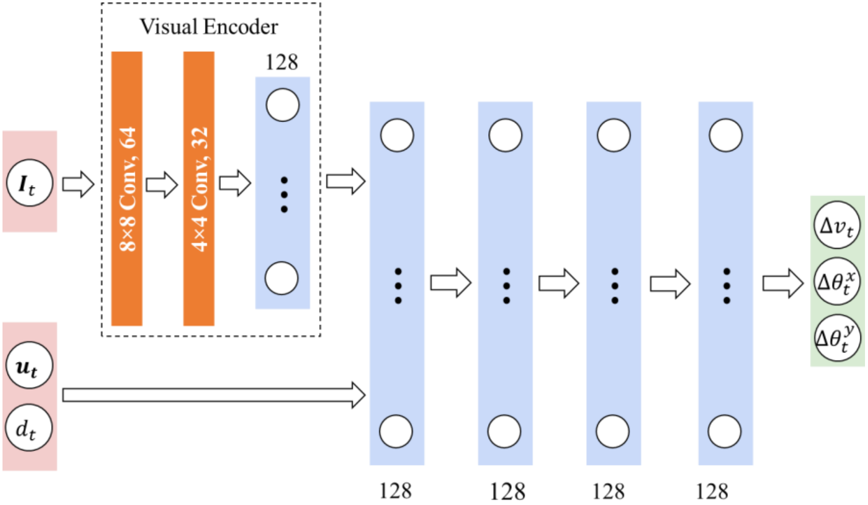

Policy network architecture

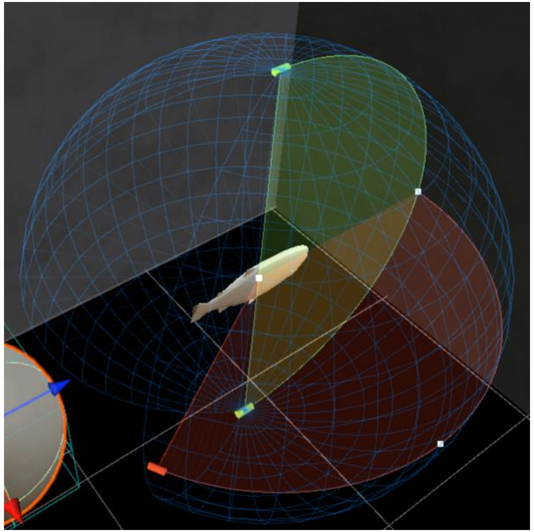

Spatial grid sensor

Experiment Setting

Species	Fish Number	Body Scale	Cage Size	Cage Shape
Coho Salmon	1000	[0.9, 1.1]	Edge of 3, height of 6	Octagon
Yellowtail	45	[0.9, 1.1]	3 x 3 x 3	Cube
Red Seabream	10	[0.9, 1.1]	3 x 3 x 3	Cube
Red Seabream	10	[0.9, 1.1]	5 x 5 x 5	Cube
Red Seabream	50	[0.9, 1.1]	3 x 3 x 3	Cube
Coho Salmon	300	0.5	Edge of 3, height of 4.6	Octagon
Coho Salmon	300	1.0	Edge of 3, height of 4.6	Octagon
Coho Salmon	300	1.5	Edge of 3, height of 4.6	Octagon
Coho Salmon	1000	0.5, 1.0 or 1.5	Edge of 1.8, height of 3	Octagon
Coho Salmon	1000	0.5, 1.0 or 1.5	Edge of 3, height of 4.6	Octagon
All three types	300	[0.9, 1.1]	Edge of 3, height of 4.6	Octagon

Eleven simulation environments

Fish Counting Model

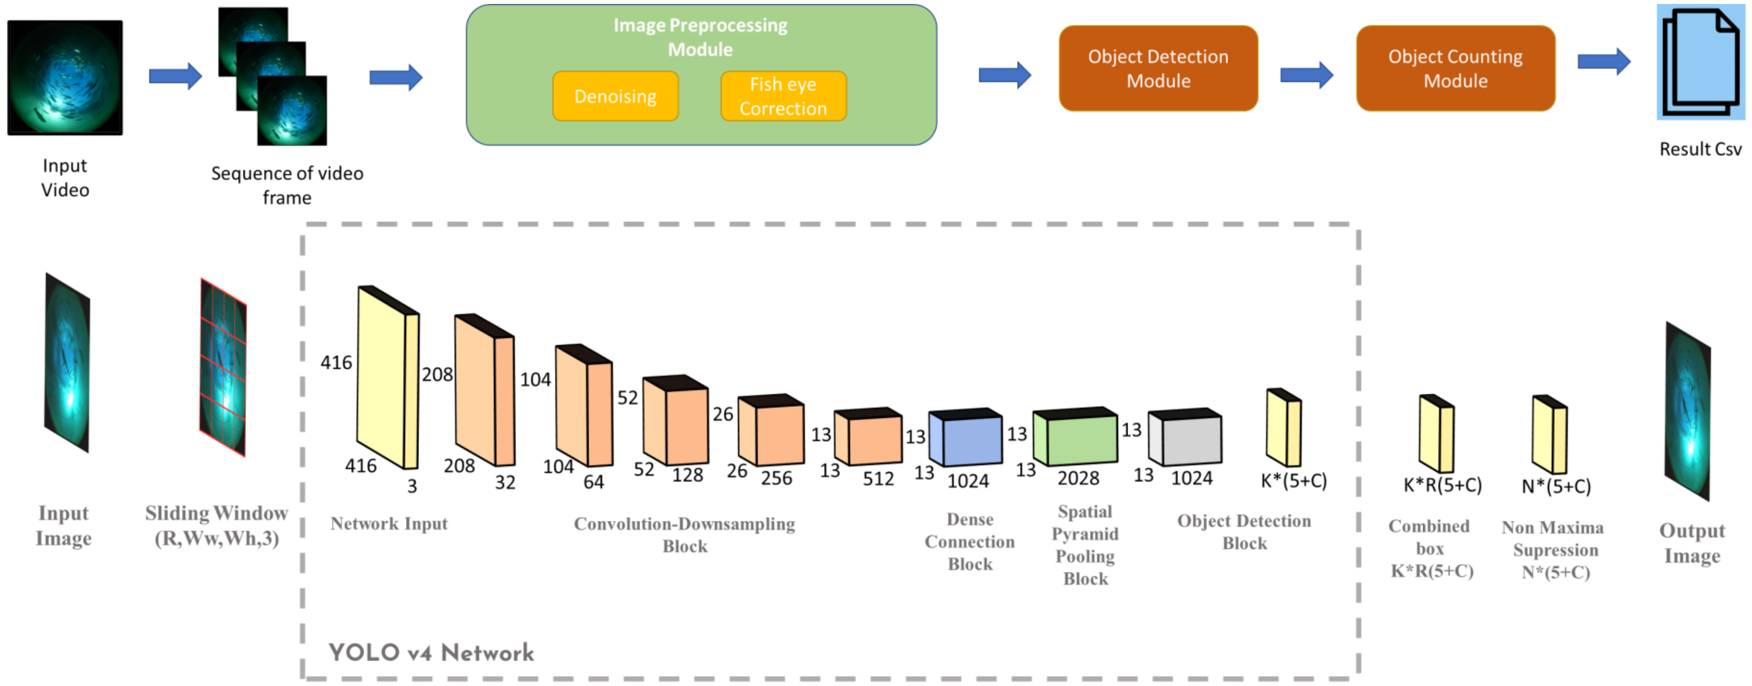

Result – Learning Curves

Coho Salmon - Pretrain

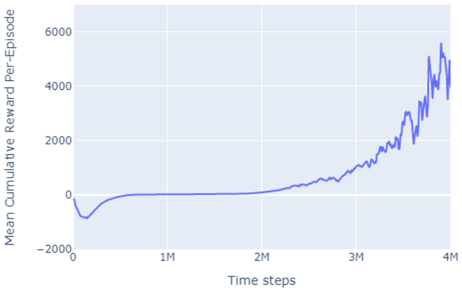

Red Seabream - Pretrain

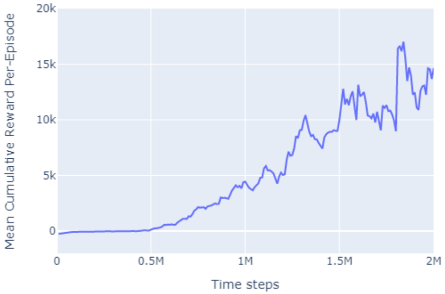

Yellowtail Amberjack - Pretrain

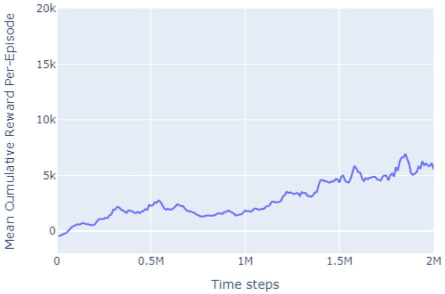

Coho Salmon - Transfer Learning

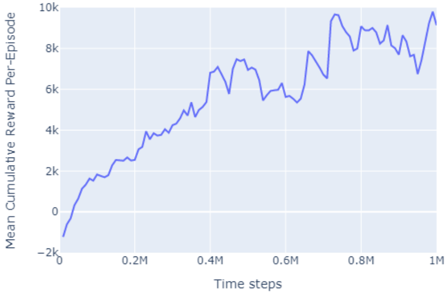

Red Seabream - Transfer Learning

Yellowtail Amberjack - Transfer Learning

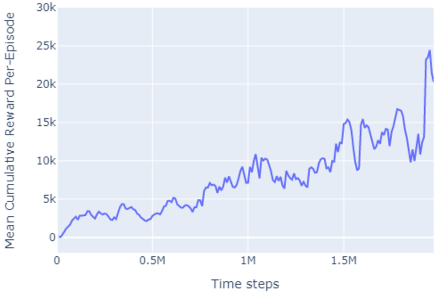

UCLA

Result – Sim vs Real

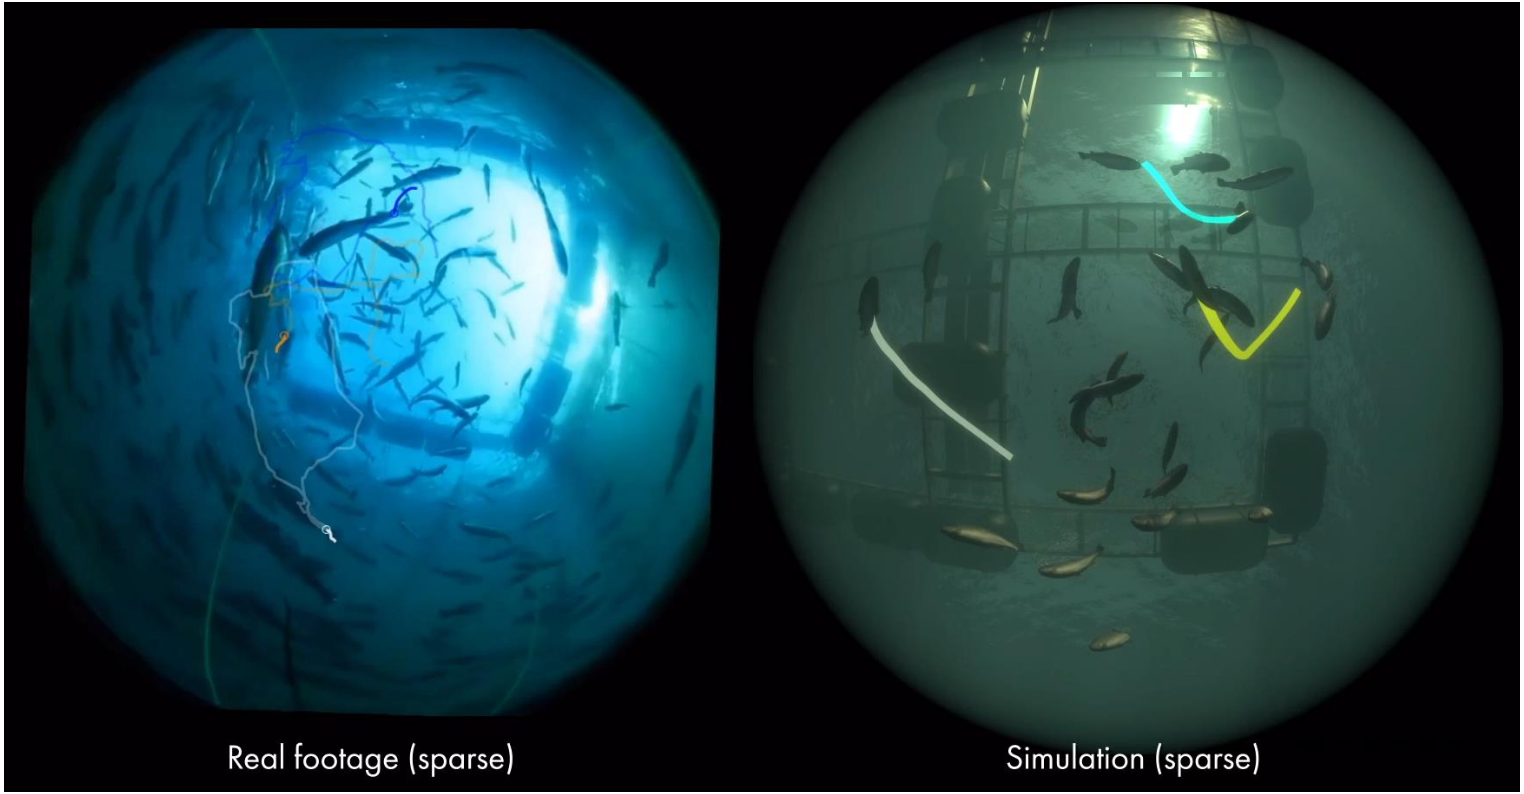

Result – Sim vs Real

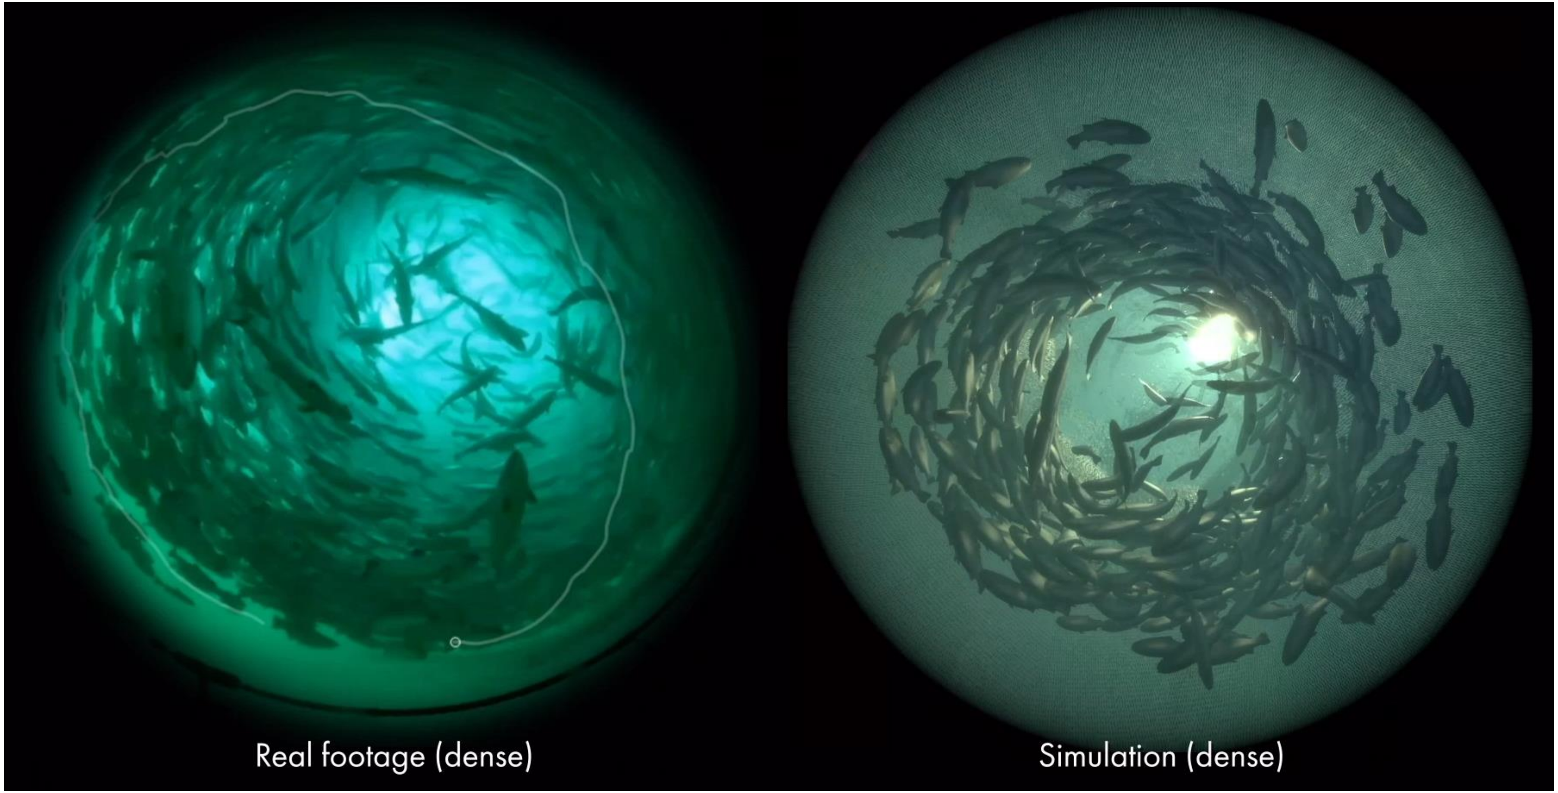

Result – Sim vs Real

Real data:

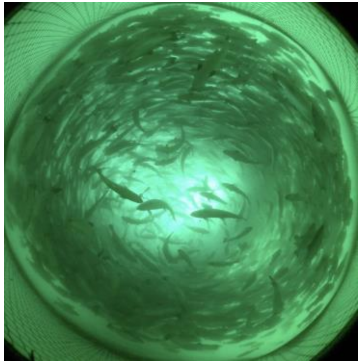

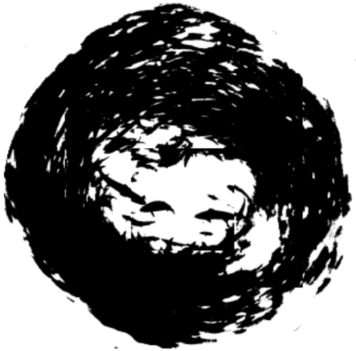

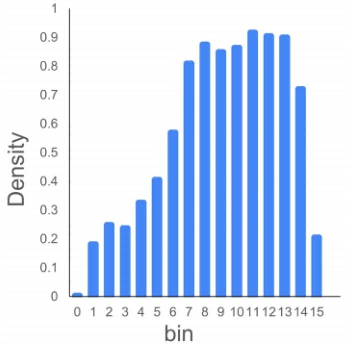

Simulation:

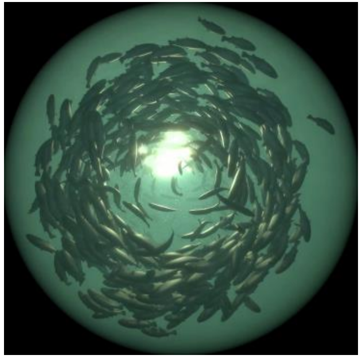

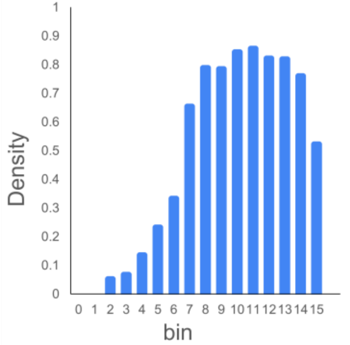

Bottom-view scene

Black-white image

Density of black pixels

Result – PPO vs SAC

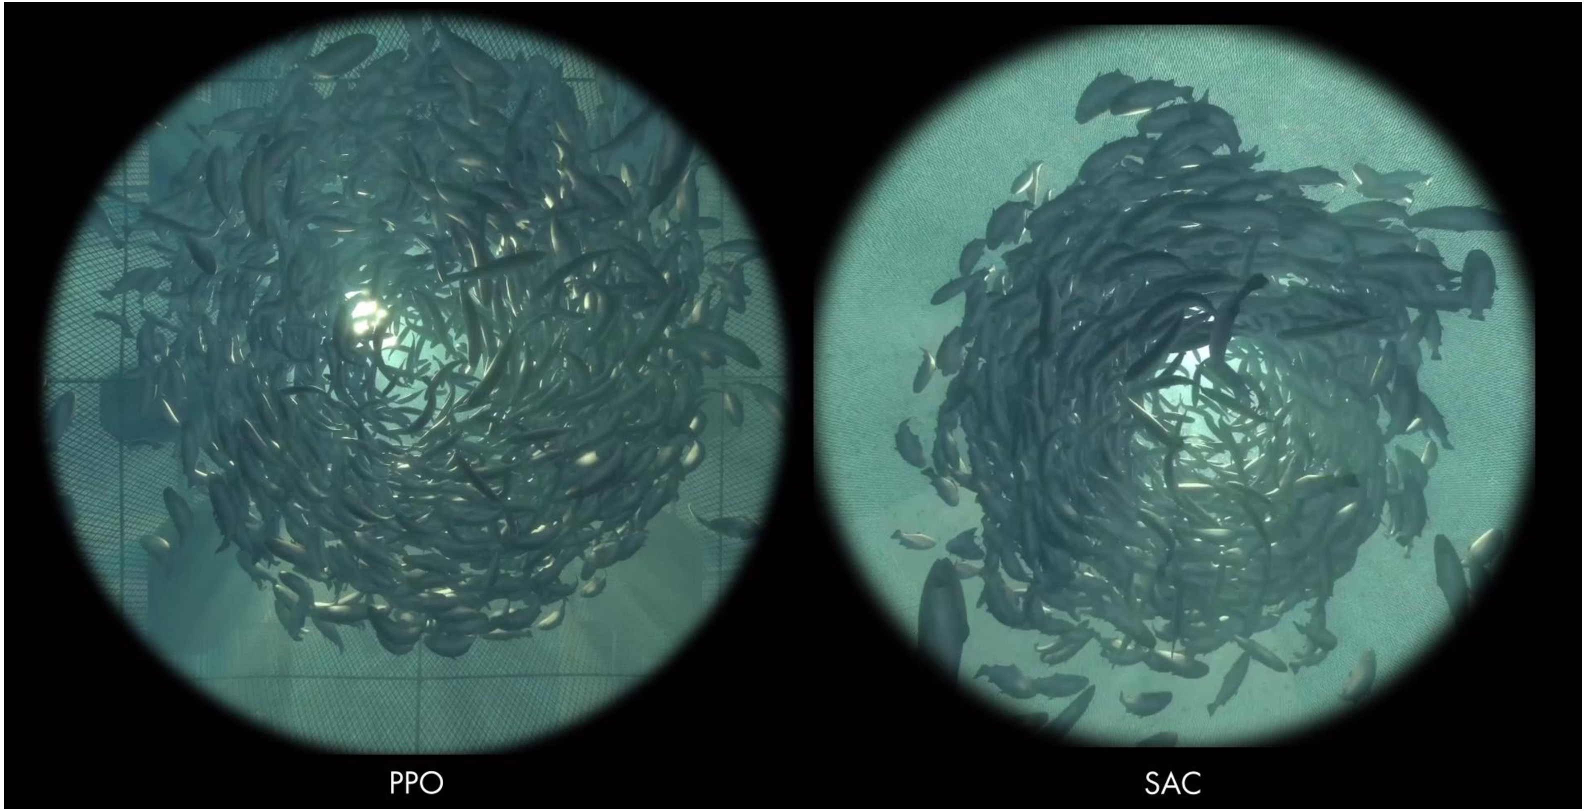

UCLA

Learning Curves – PPO vs SAC

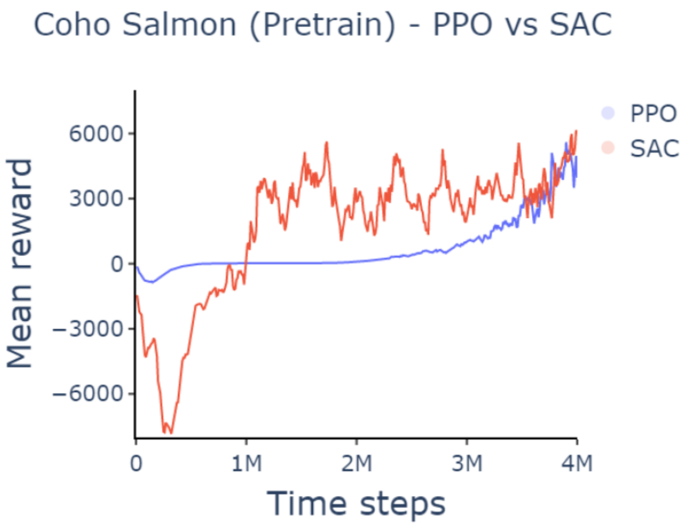

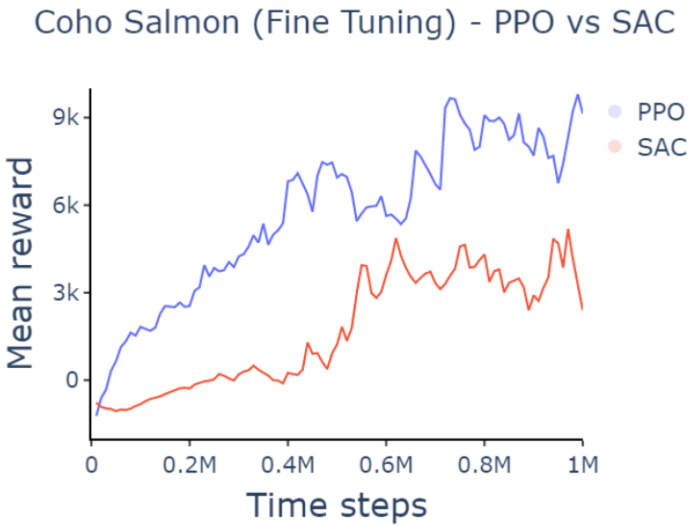

Result – Social Ranking

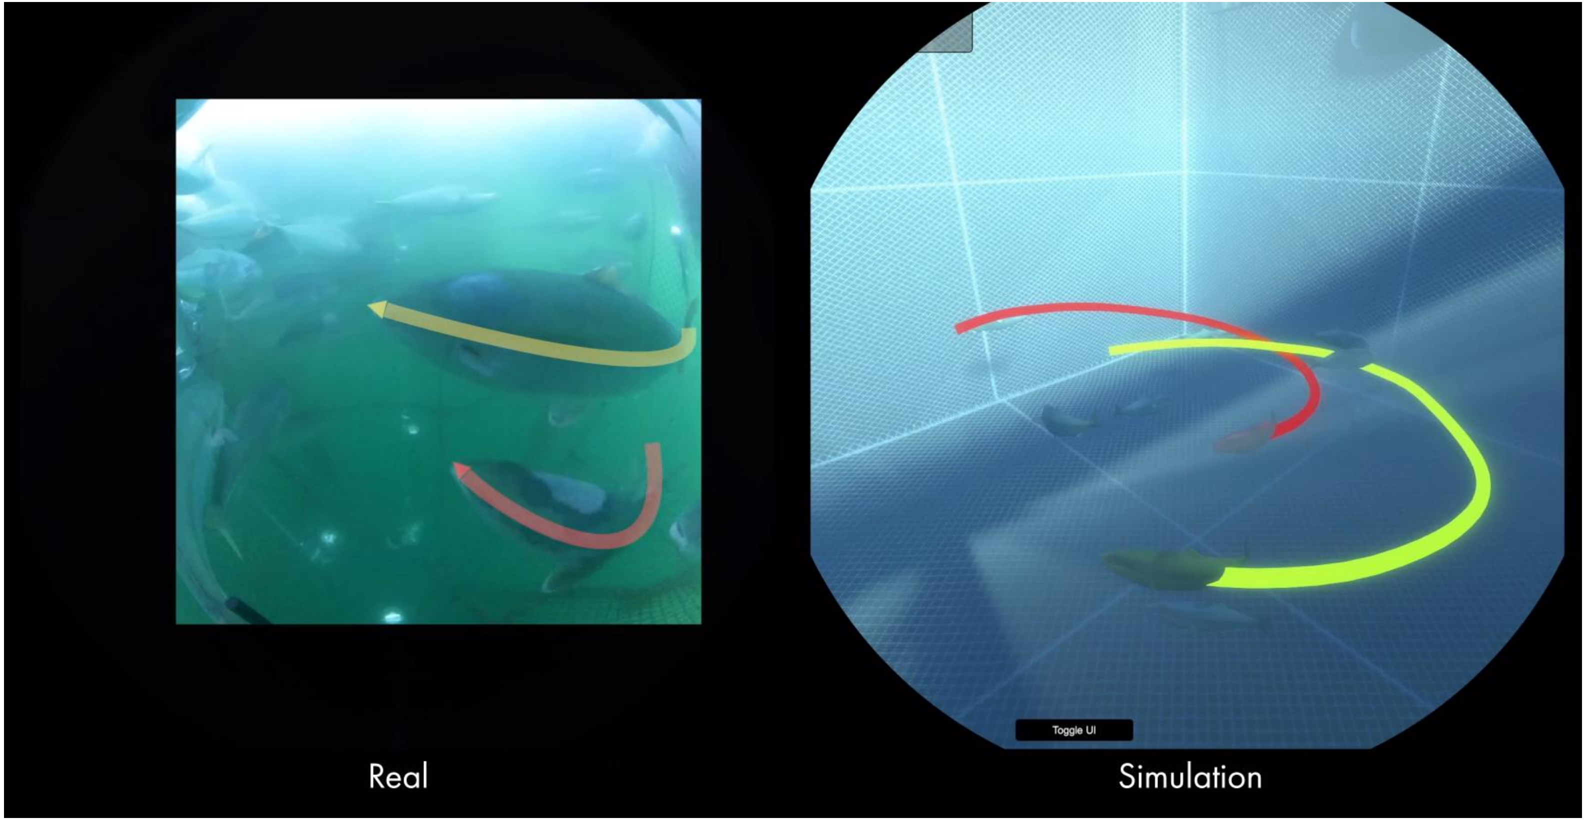

Result – Synthesized Dataset

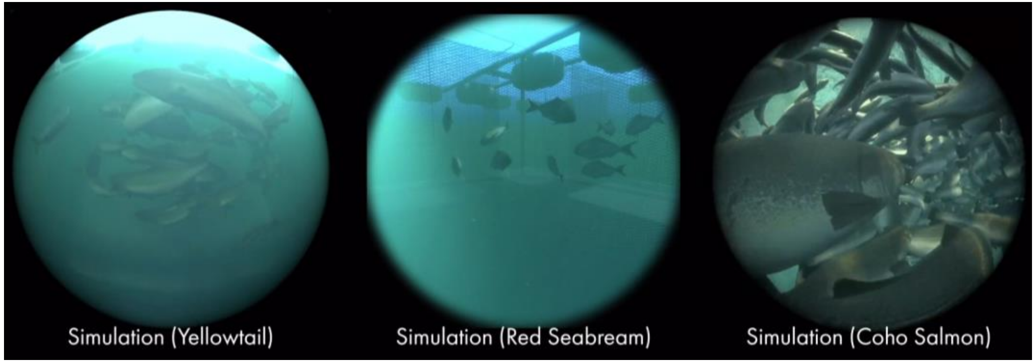

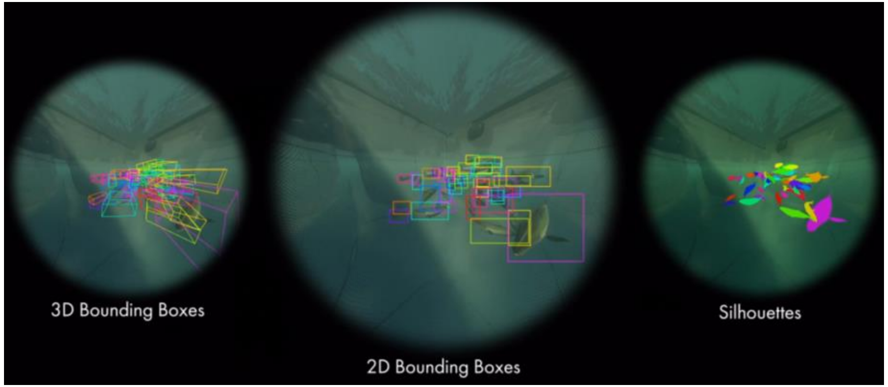

Result – Trained CV Model

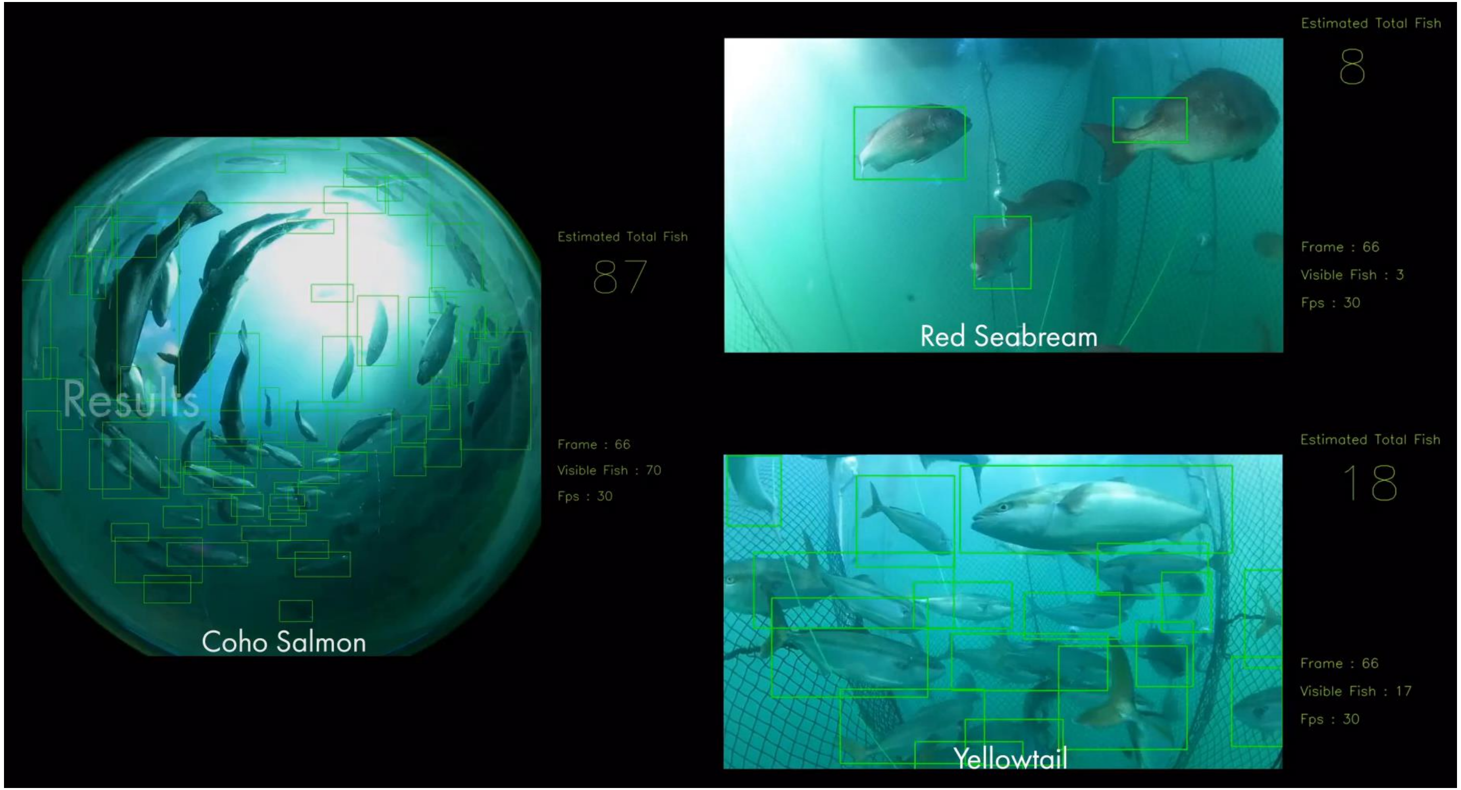

Summary

- The first biomechanical face-head-neck animation system that is capable of learning to reproduce expressions and head orientations through neuromuscular control
- A deep neuromuscular controller learns to map between FACS Action Units extracted form human facial videos and the activations of the muscle actuator that drive the biomechanical system
- An automated processing pipeline for animating expressions and head poses that can potentially be applied on any muscle-driven model

Summary

- A bio-inspired fish simulation system with DRL which allows the system to adapt to various conditions autonomously by applying transfer learning
- Together with the physically-based underwater environment simulation, the adaptive fish behavioral model can simulate a broad range of fish behaviors in many scenarios
- The proposed system is able to synthesize a large dataset with accurate annotations that can be used for computer vision tasks while avoiding the human effort of data capturing or labeling

Future Work

For the face-head-neck model:

- Add more muscles to activate facial AUs for a higher anatomical accuracy
- Extend the system to explicitly control for the lisp and eyes
- Similar control methodology can potentially be used to control a musculoskeletal hand-forearm complex

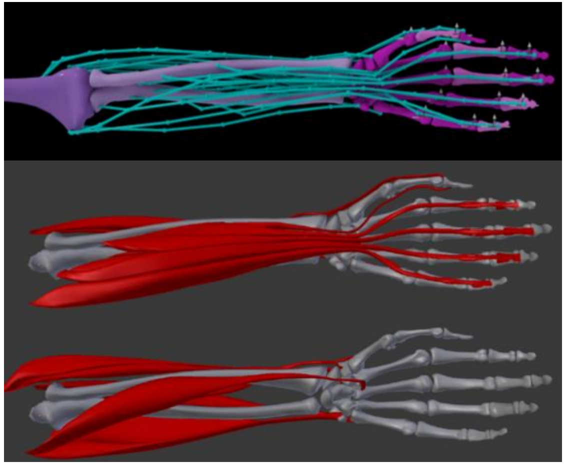

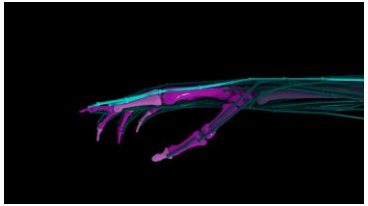

Finger extension

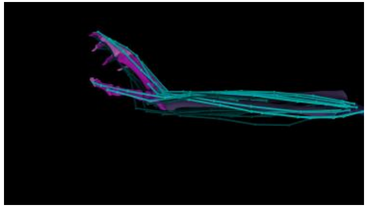

Wrist extension

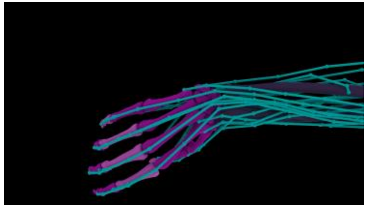

Wrist abduction

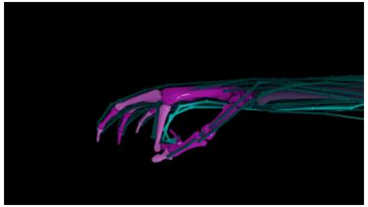

Finger flexion

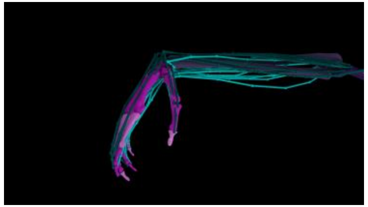

Wrist flexion

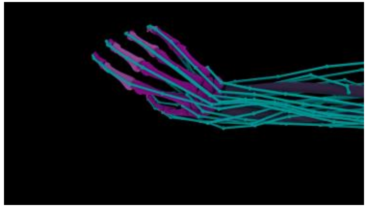

Wrist abduction

UCLA

Future Work

For the DeepFoids model:

- Extend the application to more computer vision tasks
- Dynamically adjust weights based on the performance of the DRL controller
- Include more factors that influence fish behavior
- Apply the system to other fish species

References

Zeng, Xiao S., Ishiwaka, Yuko, et al. (2022). “DeepFoids: Adaptive Bio-Inspired Fish Simulation with Deep Reinforcement Learning.” Advances in Neural Information Processing System (NeurIPS) 2022.

Ishiwaka, Yuko, Zeng, Xiao S., et al. (2021). “Foids: Bio-Inspired Fish Simulation for Generating Synthetic Datasets.” ACM Transactions on Graphics (TOG), 40(6), 1–15.

Lee, Sung-Hee, and Demetri Terzopoulos. "Heads up! Biomechanical modeling and neuromuscular control of the neck." ACM SIGGRAPH 2006 Papers. 2006. 1188-1198.

Si, Weiguang. Realistic simulation and control of human swimming and underwater movement. University of California, Los Angeles, 2013.

Zeng, Xiao S., et al. "Neuromuscular Control of the Face-Head-Neck Biomechanical Complex with Learning-Based Expression Transfer from Images and Videos." International Symposium on Visual Computing. Springer, Cham, 2021.

Thank You