Proposals
Prepare & Submit Proposals
Demo Site: Prepare Proposals
Proposal Status
Reviews & Meetings
Reviews, Panels, and Other Meetings
Reviewing for the National Science Foundation
Awards & Reporting
Project Reports
Demo Site: Project Reports
Notifications & Requests
Award Documents
Supplemental Funding Requests
Demo Site: Supplemental Funding Requests
Continuing Grant Increments Reports
PAR Research Products
Fellowships & Opportunities
Graduate Research Fellowship Program (GRFP) (Applicants, Fellows)
Graduate Research Fellowship Program (GRFP) (Officials)
Manage Reference Letters (GRFP and Postdoc Writers)
Education and Training Application (ETAP)
Manage Financials
ACM$ (Award Cash Management $ervice)
Program Income Reporting
Foreign Financial Disclosure Report (FFDR)
Individual Banking
Grant Post-Award Financial Contacts
Administration
User Management

Preview of Award 2228092 - Annual Project Report

Cover |
Accomplishments |
Products |
Participants/Organizations |
Impacts |
Changes/Problems
| Special Requirements

Cover

Federal Agency and Organization Element to Which Report is Submitted:	4900
Federal Award or Other Identifying Number Assigned by Agency:	2228092
Project Title:	CPS: Small: Learning How to Control: A Meta-Learning Approach for the Adaptive Control of Cyber-Physical Systems
PD/PI Name:	Michael D Lemmon, Principal Investigator
Recipient Organization:	University of Notre Dame
Project/Grant Period:	06/15/2023 - 05/31/2026
Reporting Period:	06/01/2024 - 05/31/2025
Submitting Official (if other than PD\PI):	N/A
Submission Date:	N/A
Signature of Submitting Official (signature shall be submitted in accordance with agency specific instructions)	N/A

Accomplishments

*** What are the major goals of the project?**

Project Overview: This project is developing methods that "learn how to control" complex cyber-physical systems (CPS) found in Manufacturing 4.0 applications. In this project, the cyber fabric is formed from a network of "digital twins" for jobshop machines and the physical fabric is formed from the actual machines on the factory floor. The project uses a meta-learning framework involving the adaptive coordination of local and global control policies. A novel adaptive model-following control strategy is used to safely adapt the controller for a machine's digital twin to the physical machine on the factory floor. A federated learning framework has been developed to learn action policies that coordinate local controllers in a manner that maximizes global jobshop performance. The project intends to evaluate the approach on a multi-robotic testbed emulating the physical flow of material through the factory.

Major goals for the project are itemized below:

[GOAL 1:] Develop learning-based approach for adapting low-level controllers of digital twins to shifts in their local environment.

[GOAL 2:] Develop learning-based distributed algorithms for global coordination of jobshop digital twins.

[GOAL 3:] Develop methods for the "safe" transfer of action policies obtained from historical datasets and digital twin models to the physical machines used on the factory floor.

[GOAL 4:] Build a multi-robotic testbed emulating the use of digital twins in managing the flow of materials across a factory floor. Demonstrate and evaluate the proposed "learning how to control" approach on this testbed.

*** What was accomplished under these goals and objectives (you must provide information for at least one of the 4 categories below)?**

Major Activities:

[Activity 1 - Learning-based control of digital twins}: This activity completed the initial model-following control architecture discussed in the project proposal. The main outcome was the hardware implementation of the learning-based model-following controller on a mini-cheetah quadruped robot. This work successfully demonstrated that the approach could safely adapt the quadruped's gait to unexpected shifts in the robot's dynamics. In particular, the base controller was adapted to stabilize the robot's gait when a box of sloshing water was added to it. Those results were reported in the 6th annual learning for dynamics & control conference (L4DC), July 15-17, 2024, Oxford, England. This activity supported GOAL 1 by demonstrating that the proposed model-following approach could "safely" adapt previously learned control policies to abrupt changes in the agent's dynamics.

[Activity 2 - Federated Learning for Global Coordination of Digital Twins:] This activity studied how federated learning can be used to control a network of digital twins in a distributed manner. Federated learning is a distributed machine learning paradigm in which multiple devices (agents) collaboratively train a global model without sharing their raw data. This project uses federated learning to learn how to coordinate the actions of a jobshop's digital twins so a global measure of jobshop performance is maximized. This year saw two major outcomes from this activity that are itemized below. This work supported GOAL 2 by developing distributed federated learning algorithms that we intend to use for the global coordination of jobshop digital twins. This work is being done by the PI and a 3rd year PhD student (Yuying Duan).

a) This activity developed a post-processing framework for evaluating the cost of enforcing local client optimality while also enforcing a global metric on the group of clients. The original work was done with respect to "algorithmic fairness" (i.e. ensuring model optimality is statistically independent of a chosen sensitive problem attribute). This work resulted in two ArXiv preprints and a paper to be presented at the 28th International Conference on Artifical Intelligence and Statistics (AISTATS), May 3-5 2025, Mai Kao Thailand (acceptance rate 30%). While the original baselines used in the AISTATS were drawn from medical network applications, future activities will shift to baselines looking at algorithmic fairness in manufacturing networks.

b) The second outcome of this activity was the development of algorithms for vertical federated learning. Vertical federated learning refers to applications where the clients have datasets drawn from different metric-measure spaces, so that labeled samples for one client cannot be directly associated with unlabeled sample in another client. Vertical federated learning develops algorithms that allow one to couple unlabeled data samples from one client with labeled data sample in another client. This problem is therefore concerned with the paucity of labeled data in real-life applications. The major outcome of this activity was the development of distributed algorithms based on the Gromov-Wasserstein (GW) transport metric that are more efficient than other examples found in the open literature. Preliminary results with our distributed GW transport algorithms are extremely promising. We are currently testing the proposed algorithms on the earlier medical baseline datasets we used for the algorithmic fairness work. Future work will apply this algorithm to the jobshop problems at the heart of this project.

[Activity 3 - Safe Sim2Real Transfer:] Digital twin action policies can be obtained through system identification based on historical data, followed by using federated learning in simulation environments to globally coordinate the digital twins. This activity is concerned with safely transferring global action policies for the digital twins to real factory floor conditions. This is sometimes referred to as Sim2Real transfer. Summer 2024 examined an approach to safe multi-agent reinforcement learning (MARL) based on finite abstractions to enforce safety. That approach, unfortunately, did not scale up to the larger scenarios expected in jobshop applications. In Spring 2025, the project began looking at Sim2Real transfer algorithms based on adaptively switching between "safe"/suboptimal local policies and "optimal" / potentially unsafe global policies. The approach uses first-step analysis of Markov chains to estimate the probability of unsafe outcomes and then using that estimate as a threshold for switching between the two policies. The work is being performed in collaboration with a soft-robotics colleague (Prof. Ozkan-Aydin) and her students. Summer 2025 activities will continue this work. This activity supports GOAL 3.

[Activity 4 - Testbed Development:] The project began supporting a new PhD students (Rafieisangari) in August 2024 to assist with testbed development. Unfortunately, the student was dismissed from the project in December 2024 for inadequate academic progress, so testbed development has been postponed until Summer 2025. The PI has prior experience with ROS programming and will begin working on the testbed in Summer 2025. Once initial development has been completed, the project will begin recruiting undergraduate and graduate assistants to help with the work. This activity supports GOAL 4.

Specific Objectives:

[Objective 1:] Develop learning-based approach for adapting low-level controllers of digital twins to shifts in their local environment (Goal 1).
Status: The model-following approach was successfully demonstrated on a quadruped robot. This year's objective is to apply that approach to jobshop task dispatching.

[Objective 2:] Develop learning-based distributed algorithms for global coordination of jobshop digital twins (Goal 2).
Status: Preliminary work in the first year of the project has demonstrated that federated learning can be used to balance competing local and global objectives in distributed decision systems. This year will apply our federated learning algorithms to the coordinated control of jobshop digital twins gathering data in different measure-metric spaces.

[Objective 3:] Develop methods for the “safe” transfer of digital twin action policies to the physical factory floor (Goal 3).
Status: Preliminary work in Spring 2025 demonstrated Sim2Real transfer by adaptively switching between safe local policies and optimal global policies. Future work will

continue to study this approach through collaborations with NDEE partners and their students.

[Objective 4:] Develop hardware multi-robot testbed (Goal 4):
Status: hardward purchased, initial robot mapping functions verified. It has, unfortunately, been difficult to find first year graduate students capable of working with the platform. As the PI has prior experience with multi-robot ROS platforms, he will develop the main parts of the testbed over Summer 2025, develop training materials, and recruit existing undergraduate and MS level graduate students into the project in the Fall.

Significant Results:

Key outcomes or Other achievements:

*** What opportunities for training and professional development has the project provided?**

The project has provided financial support for Notre Dame PhD students (Roghayeh Rafieisangari, Yuying Duan). The project has also provided opportunities for collaborating with PhD students (Nnamdi Chikere, Azhang Nazaripouya, Gelei Xu) who are financially supported by other related programs.

*** Have the results been disseminated to communities of interest? If so, please provide details.**

Results have been disseminated through conference and arXiv papers, as well as websites

Conference Papers:

1) A learning-based framework to adapt legged robots on-the-fly to unexpected disturbances, N. Fey, H. Li, N. Adrian, P. Wensing, and M.D. Lemmon, 6th Annual Learning for Dynamic and Control Conference (L4DC), July 15-17, 2024, Oxford England

2) The cost of local and global fairness in Federated Learning, Y. Duan, G. Xu, Y. Shi, and M.D. Lemmon, International Conference on Artificial Intelligence and Statistics (AISTATS), May 3-5, 2025, Mai Kao Thailand

arXiv Preprints:

1) The cost of local and lgobal fairness in federated learning (full AISTATS version), Y. Duan, G. Xu, Y. Shi, M.D. Lemmon, arXiv 2503,22762, 2025

2) Post-fair Federated Learning: achieving group and community fairness in federated learning via post-processing (preprint), Y. Duan, Y Tian, N. Chawla, M.D. Lemmon, arXiV:2405.17782, 2024

*** What do you plan to do during the next reporting period to accomplish the goals?**

1) Continue working with current PhD student (Yuying Duan) to completed work on federated learning with regard to algorithmic fairness, domain adaptation, vertifical federated learning, and multi-agent reinforcement learning. Ms. Duan should be taking her PhD Candidacy exam this next year.

2) Continue working with ND robotics colleague (Prof. Yasemin Ozka-Aydin) regarding the use of distributed learning and Sim2Real transfer for bio-inspired jobshop scheduling

3) Complete development of robotic testbed for a single robot over the summer, write training materials for the testbed, recruit undergraduate and graduate students to assist with final testbed development.

Supporting Files

Filename	Description	Uploaded By	Uploaded On
AIstats_camera_ready.pdf	The cost of local and global fairness in Federated Learning, Y. Duan, G. Xu, Y. Shi, and M.D. Lemmon, International Conference on Artificial Intelligence and Statistics (AISTATS), May 3-5, 2025, Mai Kao Thailand	Michael Lemmon	05/03/2025

Products

Books

Book Chapters

Inventions

Journals or Juried Conference Papers

View all journal publications currently available in the NSF Public Access Repository for this award.

The results in the NSF Public Access Repository will include a comprehensive listing of all journal publications recorded to date that are associated with this award.

N. Fey, H. Li, N. Adrian, P. Wensing, and M.D. Lemmon "A learning-based framework to adapt legged robots on-the-fly to unexpected disturbances" To appear in 6th annual learning for dynamics & control conference (L4DC) July 15-17, 2024, Oxford, England. Nolan Fey will present the paper at Oxford in July. Final version of the paper will be published electronically in the Proceedings of Machine Learning Research (PMLR) and will appear on project's website (https://www3.nd.edu/~lemmon/Projects/NSF-21-551/). Status = PUBLISHED.

Licenses

Other Conference Presentations / Papers

Y. Duan, G. Xu, Y. Shi, M.D. Lemmon (2025). *The cost of local and global fairness in Federated Learning*. arXiv. . Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Other Products

Other Publications

Yuying Duan, Yijun Tian, Nitesh Chawla, Michael Lemmon (2024). *Post-Fair Federated Learning: achieving group and community fairness in Federated Learning via Post-processing*. arXiv preprint arXiv.2405.17782, 2024 Federated Learning (FL) is a distributed machine learning framework in which a set of local communities collaboratively learn a shared global model while retaining all training data locally within each community. Two notions of fairness have recently emerged as important issues for federated learning: group fairness and community fairness. Group fairness requires that a model's decisions do not favor any particular group based on a set of legally protected attributes such as race or gender. Community fairness requires that global models exhibit similar levels of performance (accuracy) across all collaborating communities. Both fairness concepts can coexist within an FL framework, but the existing literature has focused on either one concept or the other. This paper proposes and analyzes a post-processing fair federated learning (FFL) framework called post-FFL. Post-FFL uses a linear program to simultaneously enforce group and community fairness while maximizing the utility of the global model. Because Post-FFL is a post-processing approach, it can be used with existing FL training pipelines whose convergence properties are well understood. This paper uses post-FFL on real-world datasets to mimic how hospital networks, for example, use federated learning to deliver community health care. Theoretical results bound the accuracy lost when post-FFL enforces both notion of fairness. Experimental results illustrate that post-FFL simultaneously improves both group and community fairness in FL. Moreover, post-FFL outperforms the existing in-processing fair federated learning in terms of improving both notions of fairness, communication efficiency and computation cost.. Status = PUBLISHED; Acknowledgement of Federal Support = Yes

Patent Applications

Technologies or Techniques

Thesis/Dissertations

Websites or Other Internet Sites

Participants/Organizations

What individuals have worked on the project?

Name	Most Senior Project Role	Nearest Person Month Worked
Lemmon, Michael	PD/PI	2
Duan, Yuying	Graduate Student (research assistant)	12
Rafieisangari, Roghayeh	Graduate Student (research assistant)	6

Full details of individuals who have worked on the project:

Michael D Lemmon
Email: lemmon@nd.edu
Most Senior Project Role: PD/PI
Nearest Person Month Worked: 2

Contribution to the Project: Principal Investigator

Funding Support: None

Change in active other support: No

International Collaboration: No
International Travel: No

Yuying Duan
Email: yduan2@nd.edu
Most Senior Project Role: Graduate Student (research assistant)
Nearest Person Month Worked: 12

Contribution to the Project: Machine Learning expert

Funding Support: 2 months - this project 10 months - departmental teaching assistant

International Collaboration: No
International Travel: No

Roghayeh Rafieisangari
Email: rrafieis@nd.edu
Most Senior Project Role: Graduate Student (research assistant)
Nearest Person Month Worked: 6

Contribution to the Project: initial technical training

Funding Support: 6 months - this project

International Collaboration: No
International Travel: No

What other organizations have been involved as partners?
Nothing to report.

Were other collaborators or contacts involved? If so, please provide details.

Nolan Fey - PhD student at MIT - for ND undergraduate

Prof. Yasemin Ozkan-Aydin - bio-inspired robotics professor in Notre Dame dept. of Electrical Engineering

Nnamdi Chikere - PhD student at ND - advisor Ozkan-Aydin

Azhang Nazaripouya - PhD student at ND - advisor Ozkan-Aydin

Impacts

What is the impact on the development of the principal discipline(s) of the project?

Project results on Sim2Real transfer are likely to provide efficient algorithmic approaches for adapting learned CPS control strategies to new test environments.

What is the impact on other disciplines?

Project results on algorithmic fairness in federated learning are likely to have a positive impact on service delivery in health care networks

Project algorithmic approaches to computing the Gromov-Wasserstein metric will solve open problems in a range of applications including domain adaptation and vertical federated learning. The algorithms will have a positive impact in addressing issues regarding the paucity of labeled data in distributed machine learning applications.

What is the impact on the development of human resources?

Project mentored two PhD students supported financially by the project. Project mentored two PhD students supported by other sources.

What was the impact on teaching and educational experiences?
Nothing to report.

What is the impact on physical resources that form infrastructure?

Project is adding additional infrastructure to Notre Dame lab facilities in the form of two deep learning workstations and the mobile robots used in developing the project's testbed.

What is the impact on institutional resources that form infrastructure?
Nothing to report.

What is the impact on information resources that form infrastructure?
Nothing to report.

What is the impact on technology transfer?
Nothing to report.

What is the impact on society beyond science and technology?
Nothing to report.

What percentage of the award's budget was spent in a foreign country?

0%

Changes/Problems

Changes in approach and reason for change
Nothing to report.

Actual or Anticipated problems or delays and actions or plans to resolve them

The first year graduate student (R. Rafieisangari) was unable to continue after her first semester due to health (anxiety) issues. She was dropped from the project after the first semester and will not return.

The issue is that the first year graduate students find it difficult to split time between required course work and robotics testbed development. To address this issue the project will have the project PI develop the testbed over the summer (he has prior ROS experience) and develop training materials for first year students. The project will then seek to encourage part-time participation of undergraduate and graduate students (MS level) in completing testbed development after having gone through the training materials.

Changes that have a significant impact on expenditures
Nothing to report.

Significant changes in use or care of human subjects
Nothing to report.

Significant changes in use or care of vertebrate animals
Nothing to report.

Significant changes in use or care of biohazards
Nothing to report.

Change in primary performance site location
Nothing to report.

Special Requirements

Responses to any special reporting requirements specified in the award terms and conditions, as well as any award specific reporting requirements.
Nothing to report.