

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.  
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 17-09-2021	2. REPORT TYPE Final Report	3. DATES COVERED (From - To) 16-Jan-2017 - 30-Apr-2020
---	--------------------------------	---

4. TITLE AND SUBTITLE Final Report: Steering T-Cell Adaptation Using Opponent Exploitation Algorithms and Computational Game Theory (Research Area 10.3)	5a. CONTRACT NUMBER W911NF-17-1-0082
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER 611102

6. AUTHORS	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES Carnegie Mellon University Associate Director, Sponsored Programs Carnegie Mellon University Pittsburgh, PA 15213 -3890	8. PERFORMING ORGANIZATION REPORT NUMBER
--	--

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211	10. SPONSOR/MONITOR'S ACRONYM(S) ARO
	11. SPONSOR/MONITOR'S REPORT NUMBER(S) 69774-CS.75

12. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.
--

13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.
---

14. ABSTRACT
--------------

15. SUBJECT TERMS
-------------------

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Tuomas Sandholm
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 412-268-8216

**RPPR Final Report**  
as of 22-Sep-2021

Agency Code: 21XD

Proposal Number: 69774CS

**Agreement Number: W911NF-17-1-0082**

**INVESTIGATOR(S):**

**Name:** Tuomas Sandholm  
**Email:** sandholm@cs.cmu.edu  
**Phone Number:** 4122688216  
**Principal:** Y

Organization: **Carnegie Mellon University**

Address: Associate Director, Sponsored Programs, Pittsburgh, PA 152133890

Country: USA

DUNS Number: 052184116

EIN: 250969449

**Report Date:** 31-Jul-2020

Date Received: 17-Sep-2021

**Final Report** for Period Beginning 16-Jan-2017 and Ending 30-Apr-2020

**Title:** Steering T-Cell Adaptation Using Opponent Exploitation Algorithms and Computational Game Theory  
(Research Area 10.3)

**Begin Performance Period:** 16-Jan-2017

**End Performance Period:** 30-Apr-2020

**Report Term:** 0-Other

Submitted By: Tuomas Sandholm

Email: sandholm@cs.cmu.edu

Phone: (412) 268-8216

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

**STEM Degrees:**

**STEM Participants:**

**Major Goals:** The objective of the project is to develop techniques for algorithmically computing sequential contingent plans for steering a T cell population to the kind of T cell population that is needed against the disease at hand (cancers versus autoimmune disease versus infections). The objective of the biological experiment portion of this project is to further develop both experimental and mathematical tools to allow for testing of novel strategies suggested by the algorithms developed in this project.

**Accomplishments:** See attached PDF.

**Training Opportunities:** Several CMU computer science PhD students have been successfully trained under this project, see the Participants section. In Prof. Morel's wet lab, several people also learned to conduct biological experiments for modeling.

**Results Dissemination:** The findings have been disseminated extensively in publications in journals, conference, and workshops.

Dr. Sandholm has also given multiple invited talks on the topic both to AI audiences and biology audiences.

Dr. Sandholm and his PhD student Gabriele Farina prepared a new course in CMU's School of Computer Science called "15-888 Computational Game Solving", which includes results from this grant. The course is offered for the first time in Fall 2021.

**Honors and Awards:** I tried to enter the awards in this box, but the box does not work correctly for large lists, even though it is supposed to support 8000 characters. SO, I am including the list of honors and awards as an uploaded PDF.

**Protocol Activity Status:**

**Technology Transfer:** Nothing to Report

**PARTICIPANTS:**

**Participant Type:** PD/PI

**RPPR Final Report**  
as of 22-Sep-2021

**Participant:** Tuomas Sandholm

**Person Months Worked:** 4.00

Project Contribution:

National Academy Member: N

**Funding Support:**

**Participant Type:** Graduate Student (research assistant)

**Participant:** Gabriele Farina

**Person Months Worked:** 15.00

Project Contribution:

National Academy Member: N

**Funding Support:**

**Participant Type:** Graduate Student (research assistant)

**Participant:** Christian Kroer

**Person Months Worked:** 14.00

Project Contribution:

National Academy Member: N

**Funding Support:**

**Participant Type:** Graduate Student (research assistant)

**Participant:** Robin Schmucker

**Person Months Worked:** 5.00

Project Contribution:

National Academy Member: N

**Funding Support:**

**International Travel:**

SWE	9 days
CAN	5 days

AUS	17 days
-----	---------

**ARTICLES:**

## RPPR Final Report as of 22-Sep-2021

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published

**Journal:** Artificial Intelligence

Publication Identifier Type:

Publication Identifier:

Volume:

Issue:

First Page #:

Date Submitted: 5/7/20 12:00AM

Date Published: 6/1/20 9:31PM

Publication Location:

**Article Title:** Limited lookahead in imperfect-information games

**Authors:** Christian Kroer, Tuomas Sandholm

**Keywords:** Game theoryEquilibrium findingLimited lookaheadImperfect-information gameNash equilibriumStackelberg equilibrium

**Abstract:** Limited lookahead has been studied for decades in perfect-information games. We initiate a new direction via two simultaneous deviation points: generalization to imperfect-information games and a game-theoretic approach. We study how one should act when facing an opponent whose lookahead is limited. We study this for opponents that differ based on their lookahead depth, based on whether they, too, have imperfect information, and based on how they break ties. We characterize the hardness of finding a Nash equilibrium or an optimal commitment strategy for either player, showing that in some of these variations the problem can be solved in polynomial time while in others it is PPAD-hard, NP-hard, or inapproximable. We proceed to design algorithms for computing optimal commitment strategies—for when the opponent breaks ties favorably, according to a fixed rule, or adversarially. We then experimentally investigate the impact of limited lookahead. The limited-lookahead player often obtains t

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published

**Journal:** Science

Publication Identifier Type:

Publication Identifier:

Volume:

Issue:

First Page #:

Date Submitted: 5/7/20 12:00AM

Date Published: 8/30/19 10:18PM

Publication Location:

**Article Title:** Superhuman AI for multiplayer poker

**Authors:** Noam Brown, Tuomas Sandholm

**Keywords:** COMP/MATH

**Abstract:** In recent years there have been great strides in artificial intelligence (AI), with games often serving as challenge problems, benchmarks, and milestones for progress. Poker has served for decades as such a challenge problem. Past successes in such benchmarks, including poker, have been limited to two-player games. However, poker in particular is traditionally played with more than two players. Multiplayer games present fundamental additional issues beyond those in two-player games, and multiplayer poker is a recognized AI milestone. In this paper we present Pluribus, an AI that we show is stronger than top human professionals in six-player no-limit Texas hold'em poker, the most popular form of poker played by humans.

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y



## RPPR Final Report as of 22-Sep-2021

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published

**Journal:** International Journal of Game Theory

Publication Identifier Type:      Publication Identifier:

Volume:      Issue:      First Page #:

Date Submitted: 10/15/20 12:00AM      Date Published:

Publication Location:

**Article Title:** Strong Nash equilibria and mixed strategies

**Authors:** Eleonora Braggion, Nicola Gatti, Roberto Lucchetti, Tuomas Sandholm, Bernhard von Stengel

**Keywords:** Noncooperative games, Strong Nash equilibrium, Mixed strategies, Pareto efficiency

**Abstract:** We study strong Nash equilibria in mixed strategies in finite games. A Nash equilibrium is strong if no coalition of players can jointly deviate so that all players in the coalition get strictly better payoffs. Our main result concerns games with two players and states that if a game admits a strong Nash equilibrium, then the payoff pairs in the support of the equilibrium lie on a straight line in the players' utility space. As a consequence, the set of games that have a strong Nash equilibrium in which at least one player plays a mixed strategy has measure zero. We show that the same property holds for games with more than two players, already when no coalition of two players can profitably deviate. Furthermore, we show that, in contrast to games with two players, in a strong Nash equilibrium an outcome that is strictly Pareto dominated may occur with positive probability.

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: **N**

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published

**Journal:** Immunology

Publication Identifier Type:      Publication Identifier:

Volume:      Issue:      First Page #:

Date Submitted: 10/16/20 12:00AM      Date Published:

Publication Location:

**Article Title:** Differential T-cell receptor signals for T helper cell programming

**Authors:** Penelope A Morel

**Keywords:** CD4 cell, T-cell receptors, regulatory T cells, signal transduction

**Abstract:** Upon encounter with their cognate antigen, naive CD4 T cells become activated and are induced to differentiate into several possible T helper (Th) cell subsets. This differentiation depends on a number of factors including antigen-presenting cells, cytokines and co-stimulatory molecules. The strength of the T-cell receptor (TCR) signal, related to the affinity of TCR for antigen and antigen dose, has emerged as a dominant factor in determining Th cell fate. Recent studies have revealed that TCR signals of high or low strength do not simply induce quantitatively different signals in the T cells, but rather qualitatively distinct pathways can be induced based on TCR signal strength. This review examines the recent literature in this area and highlights important new developments in our understanding of Th cell differentiation and TCR signal strength.

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: **Y**

# RPPR Final Report

## as of 22-Sep-2021

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 3-Accepted

**Journal:** arXiv

Publication Identifier Type: Other

Publication Identifier: arXiv:2105.12954

Volume:      Issue:

First Page #:

Date Submitted: 8/17/21 12:00AM

Date Published:

Publication Location:

**Article Title:** Better Regularization for Sequential Decision Spaces: Fast Convergence Rates for Nash, Correlated, and Team Equilibria

**Authors:** GABRIELE FARINA, CHRISTIAN KROER, TUOMAS SANDHOLM

**Keywords:** Nash, iterative first-order methods, computing equilibria

**Abstract:** We study the application of iterative first-order methods to the problem of computing equilibria of large-scale two-player extensive-form games. First-order methods must typically be instantiated with a regularizer that serves as a distance-generating function for the decision sets of the players. For the case of two-player zero-sum games, the state-of-the-art theoretical convergence rate for Nash equilibrium is achieved by using the dilated entropy function. In this paper, we introduce a new entropy-based distance-generating function for two-player zero-sum games, and show that this function achieves significantly better strong convexity properties than the dilated entropy, while maintaining the same easily-implemented closed-form proximal mapping. Extensive numerical simulations show that these superior theoretical properties translate into better numerical performance as well.

**Distribution Statement:** 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info  
**Acknowledged Federal Support:** Y

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 3-Accepted

**Journal:** arXiv

Publication Identifier Type: Other

Publication Identifier: arXiv:2002.10524

Volume:      Issue:

First Page #:

Date Submitted: 8/17/21 12:00AM

Date Published:

Publication Location:

**Article Title:** Efficient exploration of zero-sum stochastic games

**Authors:** Carlos Martin, Tuomas Sandholm

**Keywords:** zero-sum stochastic games

**Abstract:** We investigate the increasingly important and common game-solving setting where we do not have an explicit description of the game but only oracle access to it through gameplay, such as in financial or military simulations and computer games. During a limited-duration learning phase, the algorithm can control the actions of both players in order to try to learn the game and how to play it well. After that, the algorithm has to produce a strategy that has low exploitability. Our motivation is to quickly learn strategies that have low exploitability in situations where evaluating the payoffs of a queried strategy profile is costly. For the stochastic game setting, we propose using the distribution of state-action value functions induced by a belief distribution over possible environments. We compare the performance of various exploration strategies for this task, including generalizations of Thompson sampling and Bayes-UCB to this new setting.

**Distribution Statement:** 4-Distribution authorized to the Department of Defense and U.S. DoD contractors only  
**Acknowledged Federal Support:** Y

# RPPR Final Report

## as of 22-Sep-2021

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 5-Submitted

**Journal:** arXiv

Publication Identifier Type: Other

Publication Identifier: arXiv:2009.10061

Volume:      Issue:

First Page #:

Date Submitted: 8/17/21 12:00AM

Date Published:

Publication Location:

**Article Title:** Faster Algorithms for Optimal Ex-Ante Coordinated Collusive Strategies in Extensive-Form Zero-Sum Games

**Authors:** Gabriele Farina, Andrea Celli, Nicola Gatti, Tuomas Sandholm

**Keywords:** game theory, faster algorithms, zero-sum games

**Abstract:** We focus on the problem of finding an optimal strategy for a team of two players that faces an opponent in an imperfect-information zero-sum extensive-form game. Team members are not allowed to communicate during play but can coordinate before the game. In that setting, it is known that the best the team can do is sample a profile of potentially randomized strategies (one per player) from a joint (a.k.a. correlated) probability distribution at the beginning of the game. In this paper, we first provide new modeling results about computing such an optimal distribution by drawing a connection to a different literature on extensive-form correlation. Second, we provide an algorithm that computes such an optimal distribution by only using profiles where only one of the team members gets to randomize in each profile. We can also cap the number of such profiles we allow in the solution. This begets an anytime algorithm by increasing the cap.

**Distribution Statement:** 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info  
Acknowledged Federal Support: Y

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 5-Submitted

**Journal:** arXiv

Publication Identifier Type: Other

Publication Identifier: arXiv:2012.13315

Volume:      Issue:

First Page #:

Date Submitted: 8/17/21 12:00AM

Date Published:

Publication Location:

**Article Title:** Generalization in portfolio-based algorithm selection

**Authors:** Maria-Florina Balcan, Tuomas Sandholm, Ellen Vitercik

**Keywords:** AI, portfolio-based algorithm selection

**Abstract:** Portfolio-based algorithm selection has seen tremendous practical success over the past two decades. This algorithm configuration procedure works by first selecting a portfolio of diverse algorithm parameter settings, and then, on a given problem instance, using an algorithm selector to choose a parameter setting from the portfolio with strong predicted performance. Oftentimes, both the portfolio and the algorithm selector are chosen using a training set of typical problem instances from the application domain at hand. In this paper, we provide the first provable guarantees for portfolio-based algorithm selection. We analyze how large the training set should be to ensure that the resulting algorithm selector's average performance over the training set is close to its future (expected) performance.

**Distribution Statement:** 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info  
Acknowledged Federal Support: Y

## RPPR Final Report as of 22-Sep-2021

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 5-Submitted  
**Journal:** arXiv  
Publication Identifier Type: Other      Publication Identifier: arXiv:2104.01520  
Volume:      Issue:      First Page #:      Date Published:  
Date Submitted: 8/17/21 12:00AM  
Publication Location:

**Article Title:** Simple Uncoupled No-Regret Learning Dynamics for Extensive-Form Correlated Equilibrium

**Authors:** Gabriele Farina, Andrea Celli, Alberto Marchesi, Nicola Gatti

**Keywords:** Game theory, multi-agent systems

**Abstract:** The existence of simple uncoupled no-regret learning dynamics that converge to correlated equilibria in normal-form games is a celebrated result in the theory of multi-agent systems. Specifically, it has been known for more than 20 years that when all players seek to minimize their internal regret in a repeated normal-form game, the empirical frequency of play converges to a normal-form correlated equilibrium. Extensive-form games generalize normal-form games by modeling both sequential and simultaneous moves, as well as imperfect information. Because of the sequential nature and presence of private information in the game, correlation in extensive-form games possesses significantly different properties than its counterpart in normal-form games, many of which are still open research directions. Extensive-form correlated equilibrium (EFCE) has been proposed as the natural extensive-form counterpart to the classical notion of correlated equilibrium in normal-form games.

**Distribution Statement:** 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info  
**Acknowledged Federal Support:** Y

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published  
**Journal:** Mathematical Programming  
Publication Identifier Type: Other      Publication Identifier: <https://doi.org/10.1007/s10107-018-1336-7>  
Volume: 179      Issue: 1      First Page #: 385  
Date Submitted: 9/7/21 12:00AM      Date Published: 1/1/20 6:46PM  
Publication Location:

**Article Title:** Faster algorithms for extensive-form game solving via improved smoothing functions

**Authors:** Christian Kroer, Kevin Waugh, Fatma Kızılcık-Karza, Tuomas Sandholm

**Keywords:** Extensive-form game · Bilinear saddle-point problem · First-order method · Nash equilibrium · Zero-sum game

**Abstract:** Sparse iterative methods, in particular first-order methods, are known to be among the most effective in solving large-scale two-player zero-sum extensive-form games. The convergence rates of these methods depend heavily on the properties of the distance-generating function that they are based on. We investigate both the theoretical and practical performance improvement of first-order methods (FOMs) for solving extensive-form games through better design of the dilated entropy function—a class of distance-generating functions related to the domains associated with the extensive-form games. By introducing a new weighting scheme for the dilated entropy function, we develop the first distance-generating function for the strategy spaces of sequential games that has only a logarithmic dependence on the branching factor of the player.

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

**Acknowledged Federal Support:** Y

## RPPR Final Report as of 22-Sep-2021

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published  
**Journal:** Mathematical Programming Series A  
Publication Identifier Type: Other      Publication Identifier: <https://doi.org/10.1007/s10107-018-1336-7>  
Volume: 179      Issue: 1-2      First Page #: 385  
Date Submitted: 9/9/21 12:00AM      Date Published: 1/1/20 4:02PM  
Publication Location: New York, NY  
**Article Title:** Faster algorithms for extensive-form game solving via improved smoothing functions  
**Authors:** Christian Kroer, Kevin Waugh, Fatma K?l?nç-Karzan, Tuomas Sandholm  
**Keywords:** Extensive-form game · Bilinear saddle-point problem · First-order method · Nash equilibrium · Zero-sum game  
**Abstract:** Sparse iterative methods, in particular first-order methods, are known to be among the most effective in solving large-scale two-player zero-sum extensive-form games. The convergence rates of these methods depend heavily on the properties of the distance-generating function that they are based on. We investigate both the theoretical and practical performance improvement of first-order methods (FOMs) for solving extensive-form games through better design of the dilated entropy function—a class of distance-generating functions related to the domains associated with the extensiveform games. By introducing a new weighting scheme for the dilated entropy function, we develop the first distance-generating function for the strategy spaces of sequential games that has only a logarithmic dependence on the branching factor of the player.  
**Distribution Statement:** 1-Approved for public release; distribution is unlimited.  
Acknowledged Federal Support: Y

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published  
**Journal:** Science  
Publication Identifier Type: Other      Publication Identifier: [10.1126/science.aao1733](https://doi.org/10.1126/science.aao1733)  
Volume:      Issue:      First Page #:  
Date Submitted: 9/9/21 12:00AM      Date Published: 12/17/17 6:36PM  
Publication Location: Washington, D.C.  
**Article Title:** Superhuman AI for heads-up no-limit poker: Libratus beats top professionals  
**Authors:** Noam Brown, Tuomas Sandholm  
**Keywords:** AI, poker  
**Abstract:** No-limit Texas hold'em is the most popular form of poker. Despite AI successes in perfect-information games, the private information and massive game tree have made no-limit poker difficult to tackle. We present Libratus, an AI that, in a 120,000-hand competition, defeated four top human specialist professionals in heads-up no-limit Texas hold'em, the leading benchmark and long-standing challenge problem in imperfect-information game solving. Our game-theoretic approach features application independent techniques: an algorithm for computing a blueprint for the overall strategy, an algorithm that fleshes out the details of the strategy for subgames that are reached during play, and a self-improver algorithm that fixes potential weaknesses that opponents have identified in the blueprint strategy.  
**Distribution Statement:** 1-Approved for public release; distribution is unlimited.  
Acknowledged Federal Support: Y

# RPPR Final Report

## as of 22-Sep-2021

**Publication Type:** Journal Article      Peer Reviewed: Y      **Publication Status:** 1-Published

**Journal:** Management Science

Publication Identifier Type: ISSN

Publication Identifier: 0025-1909

Volume: 65

Issue: 4

First Page #: 1768

Date Submitted: 9/10/21 12:00AM

Date Published: 4/1/19 3:12PM

Publication Location: Catonsville, MD

**Article Title:** Failure-Aware Kidney Exchange

**Authors:** John Dickerson, Ariel Procaccia, Tuomas Sandholm

**Keywords:** kidney exchange • stochastic matching • stochastic set packing • maximum expected weight cycle cover • random graphs

**Abstract:** Algorithmic matches in fielded kidney exchanges do not typically result in an actual transplant. We address the problem of cycles and chains in proposed matches failing after the matching algorithm has committed to them. We show that failure-aware kidney exchange can significantly increase the expected number of lives saved (i) in theory, on random graph models; (ii) on real data from kidney exchange match runs between 2010 and 2014; and (iii) on synthetic data generated via a model of dynamic kidney exchange. This gain is robust to uncertainty over the true underlying failure rate. We design a branch-and-price-based optimal clearing algorithm specifically for the probabilistic exchange clearing problem and show that this new solver scales well on large simulated data, unlike prior clearing algorithms. Finally, we show that failure-aware matching can increase overall system efficiency and simultaneously increase the expected number of transplants to highly sensitized patients.

**Distribution Statement:** 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y

### CONFERENCE PAPERS:

**Publication Type:** Conference Paper or Presentation      **Publication Status:** 1-Published

**Conference Name:** ACM Conference on Economics and Computation

Date Received: 07-May-2020

Conference Date: 19-Jun-2018

Date Published:

Conference Location: Ithaca, NY

**Paper Title:** A General Theory of Sample Complexity for Multi-Item Profit Maximization

**Authors:** Maria-Florina Balcan, Tuomas Sandholm, Ellen Vitercik

Acknowledged Federal Support: Y

**Publication Type:** Conference Paper or Presentation      **Publication Status:** 1-Published

**Conference Name:** Association for the Advancement of Artificial Intelligence

Date Received: 07-May-2020

Conference Date: 08-Feb-2020

Date Published:

Conference Location: New York, NY

**Paper Title:** Coarse Correlation in Extensive-Form Games

**Authors:** Gabriele Farina, Tommaso Bianchi, Tuomas Sandholm

Acknowledged Federal Support: Y

**Publication Type:** Conference Paper or Presentation      **Publication Status:** 1-Published

**Conference Name:** Conference on Neural Information Processing Systems

Date Received: 07-May-2020

Conference Date: 08-Dec-2019

Date Published:

Conference Location: Vancouver, Canada

**Paper Title:** Correlation in Extensive-Form Games: Saddle-Point Formulation and Benchmarks

**Authors:** Gabriele Farina, Chun Kai Ling, Fei Fang, Tuomas Sandholm

Acknowledged Federal Support: Y

**RPPR Final Report**  
as of 22-Sep-2021

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** Conference on Neural Information Processing Systems  
Date Received: 07-May-2020 Conference Date: 03-Dec-2018 Date Published:  
Conference Location: Montreal, Canada  
**Paper Title:** Depth-Limited Solving for Imperfect-Information Games  
**Authors:** Noam Brown, Tuomas Sandholm, Brandon Amos  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** International Joint Conference on Artificial Intelligence  
Date Received: 07-May-2020 Conference Date: 11-Jul-2020 Date Published:  
Conference Location: Yokohama, Japan  
**Paper Title:** Efficient Algorithms for Learning Revenue-Maximizing Two-Part Tariffs  
**Authors:** Maria-Florina Balcan, Siddharth Prasad, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** Conference on Neural Information Processing Systems  
Date Received: 07-May-2020 Conference Date: 08-Dec-2019 Date Published:  
Conference Location: Vancouver, Canada  
**Paper Title:** Efficient Regret Minimization Algorithm for Extensive-Form Correlated Equilibrium  
**Authors:** Gabriele Farina, Chun Kai Ling, Fei Fang, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** ACM Conference on Economics and Computation  
Date Received: 07-May-2020 Conference Date: 24-Jun-2019 Date Published:  
Conference Location: Phoenix, AZ  
**Paper Title:** Estimating Approximate Incentive Compatibility  
**Authors:** Maria-Florina Balcan, Tuomas Sandholm, Ellen Vitercik  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** Conference on Neural Information Processing Systems  
Date Received: 07-May-2020 Conference Date: 03-Dec-2018 Date Published:  
Conference Location: Montreal, Canada  
**Paper Title:** Ex ante coordination and collusion in zero-sum multi-player extensive-form games  
**Authors:** Gabriele Farina, Andrea Celli, Nicola Gatti, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** International Conference on Machine Learning  
Date Received: 07-May-2020 Conference Date: 10-Jul-2018 Date Published:  
Conference Location: Stockholm, Sweden  
**Paper Title:** Learning to Branch  
**Authors:** Maria-Florina Balcan, Travis Dick, Tuomas Sandholm, Ellen Vitercik  
Acknowledged Federal Support: **Y**

**RPPR Final Report**  
as of 22-Sep-2021

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** Association for the Advancement of Artificial Intelligence  
Date Received: 07-May-2020 Conference Date: 01-Feb-2019 Date Published:  
Conference Location: Honolulu, Hawaii  
**Paper Title:** Online Convex Optimization for Sequential Decision Processes and Extensive-Form Games  
**Authors:** Gabriele Farina, Christian Kroer, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** Conference on Neural Information Processing Systems  
Date Received: 07-May-2020 Conference Date: 08-Dec-2019 Date Published:  
Conference Location: Vancouver, Canada  
**Paper Title:** Optimistic Regret Minimization for Extensive-Form Games via Dilated Distance-Generating Functions  
**Authors:** Gabriele Farina, Christian Kroer, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** Conference on Neural Information Processing Systems  
Date Received: 07-May-2020 Conference Date: 03-Dec-2018 Date Published:  
Conference Location: Montreal, Canada  
**Paper Title:** Practical Exact Algorithm for Trembling-Hand Equilibrium Refinements in Games  
**Authors:** Gabriele Farina, Nicola Gatti, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** International Conference on Machine Learning  
Date Received: 07-May-2020 Conference Date: 06-Aug-2017 Date Published:  
Conference Location: Sydney, Australia  
**Paper Title:** Reduced Space and Faster Convergence in Imperfect-Information Games via Pruning  
**Authors:** Noam Brown, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** International Conference on Machine Learning  
Date Received: 07-May-2020 Conference Date: 10-Jun-2019 Date Published:  
Conference Location: Long Beach, CA  
**Paper Title:** Regret Circuits: Composability of Regret Minimizers  
**Authors:** Gabriele Farina, Christian Kroer, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** International Conference on Machine Learning  
Date Received: 07-May-2020 Conference Date: 06-Aug-2017 Date Published:  
Conference Location: Sydney, Australia  
**Paper Title:** Regret Minimization in Behaviorally-Constrained Zero-Sum Games  
**Authors:** Gabriele Farina, Christian Kroer, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**RPPR Final Report**  
as of 22-Sep-2021

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** Association for the Advancement of Artificial Intelligence  
Date Received: 07-May-2020 Conference Date: 02-Feb-2018 Date Published:  
Conference Location: New Orleans, LA  
**Paper Title:** Robust Stackelberg Equilibria in Extensive-Form Games and extension to Limited Lookahead  
**Authors:** Christian Kroer, Gabriele Farina, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** Conference on Neural Information Processing Systems  
Date Received: 07-May-2020 Conference Date: 04-Dec-2017 Date Published:  
Conference Location: Long Beach, CA  
**Paper Title:** Safe and Nested Subgame Solving for Imperfect-Information Games  
**Authors:** Noam Brown, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** International Joint Conference on Artificial Intelligence  
Date Received: 07-May-2020 Conference Date: 19-Aug-2017 Date Published:  
Conference Location: Melbourne, Australia  
**Paper Title:** Smoothing Method for Approximate Extensive-Form Perfect Equilibrium  
**Authors:** Christian Kroer, Gabriele Farina, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** Association for the Advancement of Artificial Intelligence  
Date Received: 07-May-2020 Conference Date: 27-Jan-2019 Date Published:  
Conference Location: Honolulu, Hawaii  
**Paper Title:** Solving Imperfect-Information Games via Discounted Regret Minimization  
**Authors:** Noam Brown, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** Conference on Neural Information Processing Systems  
Date Received: 07-May-2020 Conference Date: 03-Dec-2018 Date Published:  
Conference Location: Montreal, Canada  
**Paper Title:** Solving Large Sequential Games with the Excessive Gap Technique  
**Authors:** Christian Kroer, Gabriele Farina, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** International Conference on Machine Learning  
Date Received: 07-May-2020 Conference Date: 10-Jun-2019 Date Published:  
Conference Location: Long Beach, CA  
**Paper Title:** Stable-Predictive Optimistic Counterfactual Regret Minimization  
**Authors:** Gabriele Farina, Christian Kroer, Noam Brown, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**RPPR Final Report**  
as of 22-Sep-2021

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** International Joint Conference on Artificial Intelligence  
Date Received: 07-May-2020 Conference Date: 13-Jul-2018 Date Published:  
Conference Location: Stockholm, Sweden  
**Paper Title:** Trembling-Hand Perfection in Extensive-Form Games with Commitment  
**Authors:** Gabriele Farina, Alberto Marchesi, Christian Kroer, Nicola Gatti, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** International Joint Conference on Artificial Intelligence  
Date Received: 07-May-2020 Conference Date: 19-Aug-2017 Date Published:  
Conference Location: Melbourne, Australia  
**Paper Title:** Operation Frames and Clubs in Kidney Exchange  
**Authors:** Gabriele Farina, John P. Dickerson, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** ACM Conference on Economics and Computation  
Date Received: 07-May-2020 Conference Date: 26-Jun-2017 Date Published:  
Conference Location: Cambridge, MA  
**Paper Title:** Theoretical and Practical Advances on Smoothing for Extensive-Form Games  
**Authors:** Christian Kroer, Kevin Waugh, Fatma Kilinc-Karzan, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** AAAI Conference on Artificial Intelligence  
Date Received: 15-Oct-2020 Conference Date: 07-Feb-2020 Date Published:  
Conference Location: New York, New York  
**Paper Title:** Learning to Optimize Computational Resources: Frugal Training with Generalization Guarantees  
**Authors:** Maria-Florina Balcan, Tuomas Sandholm, Ellen Vitercik  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** AAAI Workshop on Reinforcement Learning in Games  
Date Received: 15-Oct-2020 Conference Date: 08-Feb-2020 Date Published:  
Conference Location: New York, New York  
**Paper Title:** Counterfactual-Free Regret Minimization for Sequential Decision Making and Extensive-Form Games  
**Authors:** Gabriele Farina, Robin Schmucker, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** AAAI Workshop on Reinforcement Learning in Games  
Date Received: 15-Oct-2020 Conference Date: 07-Feb-2020 Date Published:  
Conference Location: New York, New York  
**Paper Title:** Composability of Regret Minimizers  
**Authors:** Gabriele Farina, Christian Kroer, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**RPPR Final Report**  
as of 22-Sep-2021

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** AAI Workshop on Reinforcement Learning in Games  
Date Received: 15-Oct-2020 Conference Date: 07-Feb-2020 Date Published:  
Conference Location: New York, New York  
**Paper Title:** Stable-Predictive Optimistic Counterfactual Regret Minimization  
**Authors:** Gabriele Farina, Christian Kroer, Noam Brown, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** Bridging Game Theory and Deep Learning Workshop NeurIPS, &#x28;Smooth Games  
Optimization and Machine Learning Series  
Date Received: 15-Oct-2020 Conference Date: 09-Dec-2019 Date Published:  
Conference Location: Vancouver  
**Paper Title:** Compositional Calculus of Regret Minimizers  
**Authors:** Gabriele Farina, Christian Kroer, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** IJCAI Workshop on Strategic Reasoning  
Date Received: 15-Oct-2020 Conference Date: 10-Aug-2019 Date Published:  
Conference Location: Macao  
**Paper Title:** Power of Correlation in Extensive-Form Games  
**Authors:** Gabriele Farina, Chun Kai Ling, Fei Fang, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** AAI Workshop on Reinforcement Learning in Games  
Date Received: 15-Oct-2020 Conference Date: 08-Feb-2020 Date Published:  
Conference Location: New York, New York  
**Paper Title:** Deep Counterfactual Regret Minimization  
**Authors:** Noam Brown, Adam Lerer, Sam Gross, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 3-Accepted  
**Conference Name:** Conference on Neural Information Processing Systems  
Date Received: 15-Oct-2020 Conference Date: 07-Dec-2020 Date Published:  
Conference Location: Virtual  
**Paper Title:** Small Nash Equilibrium Certificates in Very Large Games  
**Authors:** Brian Hu Zhang, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 3-Accepted  
**Conference Name:** Conference on Neural Information Processing Systems  
Date Received: 15-Oct-2020 Conference Date: 07-Dec-2020 Date Published:  
Conference Location: Virtual  
**Paper Title:** Polynomial-Time Computation of Optimal Correlated Equilibria in Two-Player Extensive-Form  
Games with Public Chance Moves and Beyond  
**Authors:** Gabriele Farina, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**RPPR Final Report**  
as of 22-Sep-2021

**Publication Type:** Conference Paper or Presentation **Publication Status:** 3-Accepted  
**Conference Name:** Conference on Neural Information Processing Systems  
Date Received: 15-Oct-2020 Conference Date: 07-Dec-2020 Date Published:  
Conference Location: Virtual  
**Paper Title:** Improving Policy-Constrained Kidney Exchange via Pre-Screening  
**Authors:** Duncan McElfresh, Michael Curry, Tuomas Sandholm, John Dickerson  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** International Conference on Machine Learning  
Date Received: 15-Oct-2020 Conference Date: 13-Jul-2020 Date Published:  
Conference Location: Virtual  
**Paper Title:** Stochastic Regret Minimization in Extensive-Form Games  
**Authors:** Gabriele Farina, Christian Kroer, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** International Conference on Machine Learning  
Date Received: 15-Oct-2020 Conference Date: 13-Jul-2020 Date Published:  
Conference Location: Virtual  
**Paper Title:** Sparsified Linear Programming for Zero-Sum Equilibrium Finding  
**Authors:** Brian Hu Zhang, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** International Conference on Machine Learning  
Date Received: 15-Oct-2020 Conference Date: 13-Jul-2020 Date Published:  
Conference Location: Virtual  
**Paper Title:** Refined Bounds for Algorithm Configuration: The Knife-edge of Dual Class Approximability  
**Authors:** Maria-Florina Balcan, Tuomas Sandholm, Ellen Vitercik  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** AAI Workshop on Reinforcement Learning in Games  
Date Received: 15-Oct-2020 Conference Date: 08-Feb-2020 Date Published:  
Conference Location: New York, New York  
**Paper Title:** Correlation in Extensive-Form Games: Saddle-Point Formulation and Benchmarks  
**Authors:** Gabriele Farina, Chun Kai Ling, Fei Fang, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** AAI Workshop on Reinforcement Learning in Games  
Date Received: 15-Oct-2020 Conference Date: 08-Feb-2020 Date Published:  
Conference Location: New York, New York  
**Paper Title:** Efficient Regret Minimization Algorithm for Extensive-Form Correlated Equilibrium  
**Authors:** Gabriele Farina, Chun Kai Ling, Fei Fang, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**RPPR Final Report**  
as of 22-Sep-2021

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** AAAI Workshop on Reinforcement Learning in Games  
Date Received: 15-Oct-2020 Conference Date: 07-Feb-2020 Date Published:  
Conference Location: New York, New York  
**Paper Title:** Optimistic Regret Minimization for Extensive-Form Games via Dilated Distance-Generating Functions  
**Authors:** Gabriele Farina, Christian Kroer, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** IJCAI Workshop on Strategic Reasoning  
Date Received: 15-Oct-2020 Conference Date: 11-Aug-2019 Date Published:  
Conference Location: Macao  
**Paper Title:** Optimistic Regret Minimization for Extensive-Form Games via Dilated Distance-Generating Functions  
**Authors:** Gabriele Farina, Christian Kroer, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** AAMAS Workshop on Games, Agents and Incentives Workshops  
Date Received: 15-Oct-2020 Conference Date: 13-May-2019 Date Published:  
Conference Location: Montreal  
**Paper Title:** Trembling-Hand Perfection in Stackelberg Sequential Games  
**Authors:** Gabriele Farina, Alberto Marchesi, Christian Kroer, Nicola Gatti, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** AAMAS Workshop on Games, Agents and Incentives Workshops  
Date Received: 15-Oct-2020 Conference Date: 14-May-2019 Date Published:  
Conference Location: Montreal  
**Paper Title:** Computing a Quasi-Perfect Stackelberg Equilibrium  
**Authors:** Alberto Marchesi, Gabriele Farina, Christian Kroer, Nicola Gatti, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** AAAI Workshop on Reinforcement Learning in Games  
Date Received: 15-Oct-2020 Conference Date: 28-Jan-2019 Date Published:  
Conference Location: Honolulu, Hawaii  
**Paper Title:** Depth-Limited Solving for Imperfect-Information Games  
**Authors:** Noam Brown, Tuomas Sandholm, Brandon Amos  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** AAAI-20  
Date Received: 07-Sep-2021 Conference Date: 07-Feb-2020 Date Published: 07-Feb-2020  
Conference Location: New York, NY  
**Paper Title:** Coarse Correlation in Extensive-Form Games  
**Authors:** Gabriele Farina, Tommaso Bianchi, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**RPPR Final Report**  
as of 22-Sep-2021

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** IJCAI-20  
Date Received: 07-Sep-2021 Conference Date: 07-Jan-2021 Date Published: 07-Jan-2021  
Conference Location: Virtual conference  
**Paper Title:** Efficient Algorithms for Learning Revenue-Maximizing Two-Part Tariffs  
**Authors:** Maria-Florina Balcan, Siddharth Prasad, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** AAAI-21  
Date Received: 09-Sep-2021 Conference Date: 02-Feb-2021 Date Published: 02-Feb-2021  
Conference Location: Virtual conference  
**Paper Title:** Generalization in portfolio-based algorithm selection  
**Authors:** Maria-Florina Balcan, Tuomas Sandholm, Ellen Vitercik  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** EC-21  
Date Received: 09-Sep-2021 Conference Date: 19-Jul-2021 Date Published: 19-Jul-2021  
Conference Location: Virtual conference  
**Paper Title:** Better Regularization for Sequential Decision Spaces: Fast Convergence Rates for Nash, Correlated, and Team Equilibria  
**Authors:** Gabriele Farina, Christian Kroer, Tuomas Sandholm  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** NeurIPS 2020  
Date Received: 09-Sep-2021 Conference Date: 06-Dec-2020 Date Published: 06-Dec-2020  
Conference Location: Virtual conference  
**Paper Title:** Improving Policy-Constrained Kidney Exchange via Pre-Screening  
**Authors:** Duncan McElfresh, Michael Curry, Tuomas Sandholm, John Dickerson  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** NeurIPS-20  
Date Received: 09-Sep-2021 Conference Date: 06-Dec-2020 Date Published: 06-Dec-2020  
Conference Location: Virtual conference  
**Paper Title:** No-Regret Learning Dynamics for Extensive-Form Correlated Equilibrium  
**Authors:** Andrea Celli, Alberto Marchesi, Gabriele Farina, Nicola Gatti  
Acknowledged Federal Support: **Y**

**Publication Type:** Conference Paper or Presentation **Publication Status:** 1-Published  
**Conference Name:** STOC-21  
Date Received: 09-Sep-2021 Conference Date: 21-Jun-2021 Date Published: 21-Jun-2021  
Conference Location: Virtual conference  
**Paper Title:** How Much Data Is Sufficient to Learn High-Performing Algorithms? Generalization Guarantees for Data-Driven Algorithm Design  
**Authors:** Maria-Florina Balcan, Dan DeBlasio, Travis Dick, Carl Kingsford, Tuomas Sandholm, Ellen Vitercik  
Acknowledged Federal Support: **Y**

**RPPR Final Report**  
as of 22-Sep-2021

**Partners**

,

I certify that the information in the report is complete and accurate:

Signature: Tuomas Sandholm

Signature Date: 9/17/21 10:19PM

## Honors and Awards During the Grant Period

Prof. Sandholm received the following awards during the reporting period.

- IJCAI John McCarthy Award, 2021. "For his significant research contributions to multiagent systems, computational economics, optimization and game playing, and their application in real-world settings". The IJCAI John McCarthy Award is intended to recognize established mid-career researchers, typically between fifteen to twenty-five years after obtaining their PhD, that have built up a major track record of research excellence in artificial intelligence. Nominees of the award will have made significant contributions to the research agenda in their area and will have a first-rate profile of influential research results. One award per year is given.
- AAAI Robert S. Englemore Award, 2021. "For outstanding research contributions in Artificial Intelligence, its application to electronic market places, the highly original use of AI in strategic multi-player games, and the application of AI to optimize organ exchanges". "The award is given to people who have demonstrated excellence in scholarship in AI, outstanding applications of AI, as well as extraordinary service to AAAI and the AI community."
- University Professor, 5/2020. The title of University Professor is the highest designation a faculty member can receive at Carnegie Mellon. The number of University Professors should be no more than about 10 percent of number of faculty at Carnegie Mellon University holding the rank of Professor.
- *Science* selected our work on superhuman multi-player poker as ***Science Breakthrough of the Year Runner-Up***, 2019.
- IJCAI Marvin Minsky Medal, 2019. Given for our AI Libratus, the first and only AI to beat top professionals in two-player no-limit Texas hold'em. (Second time ever this medal has been given.)
- Alumnus of the Year, Aalto University School of Science, 2019.
- Angel Jordan Professor of Computer Science, 2/2018. Endowed chair.
- Allen Newell Award for Research Excellence, Carnegie Mellon University, 2018. For our "work in advancing artificial intelligence in hidden-information adversarial settings and demonstrating this via Libratus".
- Annual Supercomputing Award for "Best use of AI", Supercomputing Conference (SC), 2018.

- La Recherche (a French science monthly very similar to Scientific American) selected Libratus as one of the scientific breakthroughs of 2017.
- *Science* selected our work (together with University of Alberta's) on "AI Mastering Poker" as ***one of the 12 runner-ups for Scientific Breakthrough of the Year*** 2017.
- Annual Supercomputing Award for "Best use of AI", Supercomputing Conference (SC), November 2017. This was for Libratus, our AI that became the first AI to beat the best humans at no-limit poker. Specifically: "By beating four of the world's best human players at heads-up, no-limit Texas hold'em poker, Libratus set a new benchmark for artificial intelligence in a game that involves incomplete information and deception."
- A similar form of poker as we used in the Brains vs. AI competition that I organized in January 2017 was selected in Fall 2017 into observer status as an Olympic sport. That announcement discussed our Brains vs. AI event specifically.
- Our work on "Computer Beats Poker Pros in Brains vs. AI" was selected as one of the *Top 10 Stories of 2017*, Pokernews.
- Libratus beating top human professionals was selected #3 in the list of Top Pittsburgh Tech Stories for 2017, by Tribune Review.
- Honorary Doctorate, University of Zurich, 2016. They also organized a symposium the day before: "Electronic Markets Design: A Symposium in Honor of Tuomas Sandholm".
- Aminer Most Influential Scholar Award, 2019.
- *100 Most Intriguing Entrepreneurs*, award from Goldman Sachs, 2020.

Dr. Sandholm's -- with his students -- best paper awards & other recognition of specific papers

- My PhD student Noam Brown's dissertation won the International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS) Victor Lesser Distinguished Dissertation award, 2021.
- My PhD student Noam Brown's dissertation won the CMU School of Computer Science Distinguished Dissertation Award, 2020. It also won CMU School of Computer Science's nomination for the ACM Doctoral Dissertation Award and CMU's nomination for the AAAI/SIGAI AI Dissertation award.
- Our paper "Superhuman AI for Multiplayer Poker" was selected to be presented in the special plenary session on *Highlights Beyond EC*, at the *ACM Conference on Economics and Computation (EC)*, 2020.

- Our paper "Solving Imperfect-Information Games with Discounted Counterfactual Regret Minimization" was selected for an *Outstanding Paper Honorable Mention* at the *Association for the Advancement of Artificial Intelligence Annual Conference (AAAI)*, 2019. (1 of 4 papers receiving recognition out of 1,150 accepted papers and 7,095 submissions.)
- Our paper "Estimating approximate incentive compatibility" won the "Exemplary Artificial Intelligence Track Paper Award" at the *ACM Conference on Economics and Computation (EC)*, 2019. Also won the Best Presentation by a Student or Postdoctoral Researcher Award. Also, invited to the *ACM Transactions on Economics and Computation (TEAC)* Special Issue on Best of EC-19.
- Our paper "Safe and Nested Subgame Solving for Imperfect-Information Games" was selected to be one of the three papers for at the inaugural annual *AGT Fest*, 2018. These are top Algorithmic Game Theory papers from other conferences and journals that get presented in a special plenary session at the *ACM Conference on Economics and Computation (EC)*.
- Our paper "Safe and Nested Subgame Solving for Imperfect-Information Games" won a **Best Paper Award** at the *Neural Information Processing Systems: Natural and Synthetic (NeurIPS)* conference, 2017. It was one of three best paper awards, out of 678 accepted papers and 3,240 submissions.
- Our paper "Faster algorithms for extensive-form game solving via improved smoothing functions" won runner-up status in INFORMS Computing Society's *2017 INFORMS Student Paper Competition*.
- Our paper "Small Representations of Big Kidney Exchange Graphs" from
- *EXPLORE-2016: The 3rd Workshop on Exploring Beyond the Worst Case in Computational Social Choice* was selected as the *Most Visionary Paper*.
- My undergraduate advisee Ben Plaut's thesis "Algorithms for Social Good: Kidney Exchange" won CMU's *2016 Allen Newell Award for Excellence in Undergraduate Research*. Ben then joined Stanford CS as a PhD student.

## Accomplished under Goals

### A. Computer Science

#### Significance

We are using algorithms to *control* – using automatically-generated sequential contingent plans – biological adaptation.

### **Accomplishments on the algorithms side:**

#### *Better integration of the old Boolean signaling pathway model*

We integrated the Boolean signaling pathway model that had been developed by immunologists (which we also used in our IJCAI-16 paper) better with our sequential planning algorithms and conducted more realistic experiments.

#### *Work on a newer non-Boolean T cell steering model*

We started experimenting on larger and more accurate models. After multiple iterations with the biologists, we converged on a continuous-time rule-based model that a student of my immunology collaborator had developed with my collaborator ([1] Sheehan, Robert. 2016. "Closing the loop: A combined computational modeling and experimental approach provides novel insights into immune cell signaling systems and their global effects". University of Pittsburgh, PhD thesis.). The following quote summarizes the model:

"we developed a novel rule-based model of the T cell differentiation pathway that is based on our Boolean model, but also considers concentrations of signaling elements and reaction rates that cannot easily be incorporated in the Boolean framework. Thus, rule-based simulations results reflect realistic time scales of protein interactions and enzyme kinetics, in addition to changes in protein and mRNA abundance, that allow tight integration between model predictions and experimental data. In addition to protein interactions, transcription, and translation, the rule based model of T cell differentiation explicitly defines post-translational modifications central to the receptor activated signaling pathways"

One of the most interesting features of the model is the modeling of a shared IL2 pool, as this allows to model a whole population of cells, each contending part of the common pool. This type of inter-cell communication was not possible in the Boolean model.

#### *New fast cell signaling pathway simulator*

We wrote an automatic system that can take any rule-based model and convert it into an efficient numerical simulator. This was important to quickly iterate with the biologists, as multiple models had to be tried and tweaked before converging to the final one. Several iterations were also needed to verify the correctness of the new simulator. As an integration test, we recreated the experimental results in [1]. This new simulator is **two orders of magnitude faster** than the old one, which was based on Bionetgen. One of the reasons for such speed increase is the fact that the simulator is now specifically tuned for one model only (as

opposed to being a generic simulator), and that we removed all the expensive communication with the Bionetgen interpreter.

### *Exploiting a limited-lookahead opponent such as a biological entity*

We got a paper accepted to the journal *Artificial Intelligence*, 2020, on this. Limited lookahead has been studied for decades in perfect-information games. We initiate a new direction via two simultaneous deviation points: generalization to imperfect-information games and a game-theoretic approach. We study how one should act when facing an opponent whose lookahead is limited. We study this for opponents that differ based on their lookahead depth, based on whether they, too, have imperfect information, and based on how they break ties. We characterize the hardness of finding a Nash equilibrium or an optimal commitment strategy for either player, showing that in some of these variations the problem can be solved in polynomial time while in others it is PPAD-hard, NP-hard, or inapproximable. We proceed to design algorithms for computing optimal commitment strategies---for when the opponent breaks ties favorably, according to a fixed rule, or adversarially. We then experimentally investigate the impact of limited lookahead. The limited-lookahead player often obtains the value of the game if she knows the expected values of nodes in the game tree for some equilibrium---but we prove this is not sufficient in general. Finally, we study the impact of noise in those estimates and different lookahead depths.

### *Robust algorithms for playing against limited-lookahead opponents*

On this work we got a paper accepted to AAI-18: "Robust Stackelberg Equilibria in Extensive-Form Games and Extension to Limited Lookahead" by Christian Kroer, Gabriele Farina, and Tuomas Sandholm. In Proceedings of the AAI Conference on Artificial Intelligence (AAI), 2018.

Here is a summary of that work:

Stackelberg equilibria have become increasingly important as a solution concept in computational game theory, largely inspired by practical problems such as security settings. In practice, however, there is typically uncertainty regarding the model about the opponent. This paper is, to our knowledge, the first to investigate Stackelberg equilibria under uncertainty in extensive-form games, one of the broadest classes of game. We introduce robust Stackelberg equilibria, where the uncertainty is about the opponent's payoffs, as well as ones where the opponent has limited lookahead and the uncertainty is about the opponent's node evaluation function. We develop a new mixed-integer program for the deterministic limited-lookahead setting. We then extend the program to the robust setting for Stackelberg equilibrium under unlimited and under limited lookahead by the opponent. We show that for the specific case of interval uncertainty about the opponent's payoffs (or about the opponent's node evaluations in

the case of limited lookahead), robust Stackelberg equilibria can be computed with a mixed-integer program that is of the same asymptotic size as that for the deterministic setting.

Our limited-lookahead results are useful for settings where it is not always desirable to model adversaries as fully rational, but as having limited lookahead capability. This includes settings such as biological games, where the goal is to steer an evolutionary process or an adaptation process which typically acts myopically without lookahead (Sandholm AAAI-15; Kroer and Sandholm IJCAI-16) (and security games where opponents are often assumed to be myopic (which can be especially well motivated when the number of adversaries is large (Yin et al. 2012) or in the case of opportunistic criminals (Zhang et al. 2016; Rosenfeld and Kraus 2017))). Our model of limited lookahead is an extension of that of Kroer and Sandholm (2015) to a robust setting. Kroer and Sandholm (2015) gave a MIP for computing an optimal strategy to commit to in the deterministic setting. We show an alternative MIP for computing such a strategy to commit to, which we then extend to the robust setting.

### *Game-theoretic algorithms against rational opponents*

As one of the proposed directions, we developed game-theoretic algorithms against rational opponents. These (for zero-sum games such as biological steering) are safe in that they cannot be exploited by any opponent. We developed the fastest algorithms for extensive-form imperfect-information games. We published a large number of papers on these, including two papers in *Science*. The techniques that we developed included better weighting of iterates, search tree pruning for imperfect-information games, sound endgame solving for imperfect-information games, and sound depth-limited midgame solving for imperfect-information games.

We also published the first scalable algorithms the main equilibrium refinement for imperfect-information games, trembling-hand refinement. Prior work only scaled to game trees with a handful of nodes. Now we can find exact trembling-hand-perfect Nash equilibria in games with hundreds of thousands of nodes.

We also published algorithms for noncooperative coalitional games.

### *Battling cancer directly*

In addition to the work on T cell steering, one of the applications of which was cancer, we also did a very large piece of work in steering cancer cells directly. The paper is on bioXiv, and it is in third-round review for a very minor revision at *PLOS Computational Biology*.

Here is a summary of that work:

The design of efficient combination therapies is a difficult key challenge in the treatment of complex diseases such as cancers. The large heterogeneity of cancers and the large number of available drugs renders exhaustive *in vivo* or even *in vitro* investigation of possible treatments impractical. In recent years, sophisticated mechanistic, ordinary differential equation-based pathways models that can predict treatment responses at a *molecular* level have been developed. However, surprisingly little effort has been put into leveraging these models to find novel therapies. In this paper we use for the first time, to our knowledge, a large-scale state-of-the-art pan-cancer signaling pathway model to identify candidates for novel combination therapies to treat individual cancer cell lines from various tissues (e.g., minimizing proliferation while keeping dosage low to avoid adverse side effects) and populations of heterogeneous cancer cell lines (e.g., minimizing the maximum or average proliferation across the cell lines while keeping dosage low). We also show how our method can be used to optimize the drug combinations used in *sequential* treatment plans---that is, optimized sequences of potentially different drug combinations---providing additional benefits. In order to solve the treatment optimization problems, we combine the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) algorithm with a significantly more scalable sampling scheme for truncated Gaussian distributions, based on a Hamiltonian Monte-Carlo method. These optimization techniques are independent of the signaling pathway model, and can thus be adapted to find treatment candidates for other complex diseases than cancers as well, as long as a suitable predictive model is available.

## B. Biology

### Approach in the biological experiments

Work on this project has focused on a robust experimental validation of the rule-based model that forms the basis of the computational framework for this project. We have focused on activation of the PI3K/Akt/mTOR signaling pathway downstream of the T cell receptor (TCR) since this pathway inhibits the induction of Treg. Specific components of the Akt/mTOR pathway are also implicated in the differentiation of specific Th populations, with mTORC1-deficient animals failing to generate Th1 and Th17 cells and mTORC2-deficient animals failing to generate Th2 cells. Akt activity depends on the phosphorylation of two sites following activation: Threonine 308 (T308) by PDK1 and Serine 473 (S473) by mTORC2. Key regulators of the pathway include PTEN, which antagonizes PI3K activity thereby inhibiting Akt activation, and mTORC2, which phosphorylates Akt on S473 (Fig. 1). These studies have allowed us to make novel observations concerning the mechanisms by which Akt and mTOR control Th cell fate and these results were recently published.

We have also further developed the rule-based model, and performed experiments to validate key concepts in the model. This work involves *in vitro* experiments aimed at tested specific

model predictions. In addition, detailed dose response and time course experiments were performed in order to develop robust parameter estimates for the model, which can then be used by our collaborators to test specific algorithms. We have also embarked on a broader analysis of the signaling pathway differences between cells destined to become Treg or Teff cells with the aim of giving us more potential elements that could be incorporated into rule-based model.

### **Significance (of what you have done)**

Regulatory T (Treg) cells play an important role in maintaining self-tolerance and preventing autoimmunity. We, and others, have shown that Foxp3-expressing, suppressive iTreg are induced when naïve CD4 T cells are exposed to low doses of Ag. iTreg expanded following low dose Ag are able to prevent autoimmune diabetes in vivo. Despite the powerful effect of Ag dose on T cell outcome precisely how this is translated into alternate signaling programs, and how these signals contribute to determining Treg versus Teff cell fate remain poorly understood. Our recent work demonstrates that the differential regulation of PTEN and mTORC2 by TCR signals of different strengths leads to different states of Akt phosphorylation being dominant at different doses. At high Ag dose both sites are phosphorylated but only the T308 site is phosphorylated when T cells are activated with low dose Ag. The different Akt states, induced by TCR signals, result in phosphorylation of nonoverlapping sets of Akt substrates, which then leads to large differences in downstream signaling and the induction of specific Th cell fates. These results suggest that different levels of TCR stimulation initiate qualitatively distinct differentiation programs and that differential regulation of alternative splicing by Akt is one of the key elements determining Th cell fate decisions. These results have been published recently and have also led to a successful grant application from the NIH, which will allow us to further explore the experimental side of this project. The identification of specific states that define a differentiation program gives us an important readout when we experiment with different methods of achieving these goals, based on the algorithms and sequences suggested by our collaborators.

### **Accomplishments (list results/ main theorems/ main experimental studies, etc, providing a copy of the paper)**

The Akt/mTOR pathway is a key driver of murine CD4<sup>+</sup> T cell differentiation, and induction of regulatory T (Treg) cells results from low TCR signal strength and low Akt/mTOR signaling. However, strong TCR signals induce high Akt activity that promotes Th cell induction. Yet, it is unclear how Akt controls alternate T cell fate decisions. We find that the strength of the TCR signal results in differential Akt enzymatic activity. Surprisingly, the Akt substrate networks associated with T cell fate decisions are qualitatively different. Proteomic profiling of Akt signaling networks during Treg versus Th induction demonstrates that Akt differentially regulates RNA processing and splicing factors to drive T cell differentiation. Interestingly,

heterogeneous nuclear ribonucleoprotein (hnRNP) L or hnRNP A1 are Akt substrates during Treg induction and have known roles in regulating the stability and splicing of key mRNAs that code for proteins in the canonical TCR signaling pathway, including CD3 $\zeta$  and CD45. Functionally, inhibition of Akt enzymatic activity results in the dysregulation of splicing during T cell differentiation, and knockdown of hnRNP L or hnRNP A1 results in the lower induction of Treg cells. Together, this work suggests that a switch in substrate specificity coupled to the phosphorylation status of Akt may lead to alternative cell fates and demonstrates that proteins involved with alternative splicing are important factors in T cell fate decisions.

This work is continuing with the generation of a novel mouse model that has a mutation in the Akt phosphorylation site of hnRNP A1. Preliminary analyses of this mouse indicate that the immune system develops normally and there are no abnormalities in the frequency or number of immune cell populations. However, we have investigated an in vivo model of Treg induction. In these studies, we have crossed our mutant mouse to the OTII TCR transgenic model, such that all of the T cells will respond to a specific peptide derived from ovalbumin. Using these cells, we have analyzed the ability of wild type (WT) and mutant OTII cells to differentiate into Treg following oral administration of ovalbumin, a classic model of oral tolerance. We find that the mutant cells show a reduced ability to induce Treg under these conditions, suggesting a defect in peripheral Treg induction. These studies will provide further elements that can be targeted in control algorithms.

#### *Further experiments validating the rule-based model of Th cell differentiation*

As we acquire new data we incorporate these into our model of Th/Treg differentiation and, further, perform experiments to validate predictions and aid in the estimation of parameters. Our model predicts that both the dose of antigen and the duration of antigen exposure affect T cell responses. To test if these two variables act in concert both were systematically varied, to test a wide range of combinations. The resulting phase diagram reveals a two-dimensional threshold between the two steady states of the system. A minimum dose and duration of stimulation are required to switch to the PTEN-low and Akt-high state. Higher antigen doses require a shorter stimulation, as the two variables combine to determine the total signaling input.

We next wanted to test if it was possible to shift this threshold. We re-ran the simulations initializing the system with half of the baseline concentration of PTEN. Reducing PTEN levels is known to have physiological effects, as PTEN heterozygosity has been shown to have implications in tumorigenesis and leads to autoimmunity and increased lymphoproliferation. Additionally, PTEN inhibitors lead to an increased expansion of Th cells and reduced expansion of Treg cells. Our model predicts that decreasing PTEN levels reduces the dose and duration of signal necessary to cross the threshold. Thus, perturbations in PTEN levels, and differences due to heterogeneity across cells could lead to different thresholds that are cell and condition

specific. We performed experiments using T cells from mice that were heterozygous for PTEN deletion, thus they expressed half the amount of PTEN as corresponding wild type cells. These experiments confirmed the model's predictions and showed that when cells express less PTEN the activation (as measured by pS6) occurs sooner and with lower antigen doses. This work is in the final stages of being prepared for publication.

We have also performed experiments on a single cell level in which we can identify those cells that have activated Akt and we measure the localization of one of the important downstream elements, the transcription factor Foxo1. Foxo1 is important for the induction of Treg as it stimulates the transcription of PTEN and Foxp3. Fully activated Akt phosphorylates Foxo1 causing it to leave the nucleus. Using Image Stream, which combines flow cytometry with microscopy, we have obtained images showing how individual cells respond to high versus low TCR signal strength. These are important studies as many of the other approaches, such as Western blot, provide an average response of a population of cells, These novel studies will allow us to refine our population-based mathematical model to incorporate the behavior of single cells. All of these data are being used to further develop and build the model that can be used for further algorithm-based approaches.