

## ACM NEWS

# Seeking Equilibria in Economics, Computer Science

By Allyn Jackson  
December 13, 2018

[Comments](#)
[PRINT](#)

|          |  |        |  |  |  |  |  |  |  |
|----------|--|--------|--|--|--|--|--|--|--|
| VIEW AS: |  | SHARE: |  |  |  |  |  |  |  |
|----------|--|--------|--|--|--|--|--|--|--|



Constantinos Daskalakis, professor of electrical engineering and computer science at the Massachusetts Institute of Technology, observed that, "In some sense, intractability of equilibria goes against their universality."

Credit: Bryce Vickmark



Credit: Nvidia

The pretty faces staring back at you above on the right might have an eerie familiarity, but they do not belong to anyone you've ever seen before—or to any human beings on earth. They are "hallucinated faces," as Constantinos Daskalakis puts it, dreamed up by a neural network trained in a machine learning framework called a Generative Adversarial Network, or GAN for short.

The concept of a GAN was first formulated in a ground-breaking 2014 [paper](#) by Ian Goodfellow et al. Since then, research about GANs has exploded, with hundreds of papers being written and top theoretical computer scientists like Daskalakis jumping into the subject.

GANs stand at the crossroads of computer science and game theory, a crossroads that Daskalakis knows well. Last summer, he received the 2018 [Rolf Nevanlinna Prize](#) from the International Mathematical Union for transforming understanding of the computational complexity of fundamental problems from economics. His work brings deep theoretical analysis to bear on high-impact problems with direct effects on real-world applications.

## Intractability of Nash equilibria

A professor in the department of electrical engineering and computer science at the Massachusetts Institute of Technology, Daskalakis first came to prominence with his Ph.D. thesis "[The Complexity of Nash Equilibria](#)," which received the 2008 [ACM Doctoral Dissertation Award](#). The starting point was John Nash's iconic 1950 [theorem](#) that revolutionized economics by showing that every game has an equilibrium; that is, a set of strategies such that no player in the game can gain by changing strategy.

## SIGN IN for Full Access

User Name

Password

» [Forgot Password?](#)

» [Create an ACM Web Account](#)

SIGN IN

## MORE NEWS & OPINIONS

### Bees With Backpacks: The Next Army of Data Collectors?

Joe Dysart

### This Is Why AI Has Yet to Reshape Most Businesses

Technology Review

### The Artificialistic Fallacy

Robin K. Hill

## ACM RESOURCES

### .NET Framework: Introduction

Courses

Nash equilibria always exist, but can players compute them? This is the question Daskalakis tackled in a 2006 paper with Paul W. Goldberg and Christos Papadimitriou.<sup>[1]</sup> They showed that the problem of computing Nash equilibria is PPAD-complete, meaning that it is at least as hard as any problem in the complexity class known as PPAD.

Invented in 1994 by Papadimitriou, PPAD stands for "polynomial parity argument for directed graphs." It contains problems whose computational complexity cannot be analyzed by appealing to the standard P versus NP paradigm. Unlike NP-complete problems, which derive their hardness from the possibility that a solution might not exist, problems in PPAD are guaranteed to have a solution, but computing the solution could be intractable.

For decades, economists had assumed that computing Nash equilibria would be a central aim of rational players in a game. By showing that such computations are intractable in general, Daskalakis's work overturned this received wisdom. "What was cool about Nash's theorem was its universality," Daskalakis said; it applies to any game, with any number of players. "But if equilibria are intractable, why would you expect that agents, who are boundedly rational and have bounded computational resources in their brains, would actually arrive at an equilibrium? In some sense, intractability of equilibria goes against their universality."

The work of Daskalakis on Nash equilibria stimulated a great deal of subsequent research. One example is the Ph.D. thesis of another Papadimitriou student, Aviad Rubinstein, who received the [2017 ACM Doctoral Dissertation Award](#) for his dissertation "[Hardness of Approximation Between P and NP](#)." Rubinstein studied *approximate* Nash equilibria, which allow for the possibility that one player might have a small incentive to change strategy. He proved that the problem of computing approximate Nash equilibria is also intractable (albeit in a weaker sense than computing exact Nash equilibria).

### Equilibria and GANs

While computing equilibria in general is intractable, there are situations in which they can be computed. Understanding those situations has taken on fresh urgency with the advent of GANs.

Deep neural networks provide a way to process complicated datasets, such as collections of natural images. The greatest successes have come in discrimination tools, in which, for example, a computer is trained to pick out pictures of cats from a set of natural images. Generative models—in which, say, a set of natural images of cats is used to train a computer to produce new pictures of cats—have been less successful.

The idea behind GANs is to exploit the discriminatory power of deep learning to improve generative power. A GAN sets up a game between two neural networks, the Generator and the Discriminator. The Generator takes random data as input and produces an object, such as an image of a cat. The Discriminator outputs a judgment about how well the generated image compares to a training set of real cat images.

This output is used to adjust the parameters of the Generator to produce images that are harder to discriminate from real images. Simultaneously, the output is used to adjust the parameters of the Discriminator to become better at discriminating between generated and real images, and the game continues.

At the Nash equilibrium of this game, two things happen: the Generator is doing the best possible job it can to fool the Discriminator, and the Discriminator is doing the best possible job it can to distinguish between real images and images produced, or "hallucinated," by the Generator.

The concept of a GAN is brilliant. Its implementation? Not always so brilliant. As Daskalakis and colleagues put it in a recent paper<sup>[2]</sup>, "GANs are very finicky to train." Part of the finickiness relates to the computational methods used to home in on equilibria, as these methods result in instabilities. Daskalakis has worked with several collaborators to come up with more robust methods, in particular by exploring variants on gradient descent.

Their methods are guaranteed to work well when the game between the Generator and the Discriminator has a convex-concave objective. They are also empirically more stable than standard methods in the non-convex-concave case, which is the most interesting case for GAN training.

Yet, the theoretical understanding is still missing. In the non-convex-concave case, "We have little idea what is happening," Daskalakis said. "But the question is mathematically fascinating, and practically important." It was one of the open questions he proposed this past September at the [Heidelberg Laureate Forum](#), an annual conference bringing young researchers together with prize-winning computer scientists and mathematicians for a week of lectures and informal interactions.

More generally, the question of how to identify families of games for which Nash equilibria are tractable is wide open. "Moving outside of the class of nicely behaved [games]," said Daskalakis, "it's a whole different ballpark for equilibrium computation algorithms."

*Allyn Jackson is a journalist specializing in science and mathematics topics, who is based in Munich, Germany.*

---

[1] The three collaborators wrote an excellent expository account, "The Complexity of Computing a Nash Equilibrium," for the February 2009 *Communications of the ACM*.

[2] "Training GANs with Optimism," version posted to the arXiv in February 2018, <https://arxiv.org/abs/1711.00141>.

---

No entries found