

AUCTION GAMES

Playing with symmetry with neural networks

Finding the optimum design of a complex auction is a challenging and important economic problem. Multi-agent deep learning can help find equilibria by making use of inherent symmetries in bidding strategies.

David C. Parkes

Auctions are used by governments worldwide to sell wireless spectrum licenses, bringing in hundreds of millions of dollars of revenue. Even with these high stakes, the optimal design of these auctions is not fully understood. A particular challenge is that their allocation and pricing rules induce situations of strategic interdependence that are hard to study analytically. Suppose two companies, one bidding on license A and one on license B, bid against a company interested in the package {A,B}. Each regional company wants to win its license and pay as little as possible while bidding enough so that their bids combine to outbid the national bidder. Auctions for this 'local–local–global' (LLG) model, simple as it is, are at the forefront of problems that can be attacked through analytical methods¹. Writing in *Nature Machine Intelligence*, Bichler et al.² illustrate an exciting advance in the use of artificial neural networks (ANNs) to find an equilibrium in these kinds of auctions. Coming with theoretical proofs of soundness, this gives a promising new direction in the use of AI for economics.

The idea is deceptively simple: use an ANN to represent a bidding policy and through self-play and gradient dynamics find an equilibrium point where bidders play a best-response to each other. And yet, learning dynamics for games are notoriously difficult to understand and worse, often fail to converge³. A bidder's strategy in an auction is a function from value to bid ('if my value is US\$1 million then I will bid US\$900,000'). It is these functions that are the subject of an equilibrium analysis. Strategies form an equilibrium when each bidder is simultaneously maximizing its own expected payoff given a known distribution on values and thus bids. Bichler et al. exploit a useful property of some auction games, namely symmetry.

Symmetry arises when the roles of bidders are exchangeable: each bidder's value comes from the same distribution, and each bidder has the same utility function, that is, the same appetite for risk. Since symmetric equilibria exist in symmetric games of this form, a second kind of symmetry can then be imposed by insisting that each bidder adopts the same strategy. These two kinds of symmetry provide enough structure for gradient dynamics to converge through an appeal to recent theoretical advances on the convergence of gradient-based dynamics in potential games⁴.

There remain some technical challenges. First, gradient-based dynamics converge only to a local equilibrium, for which the guarantee is that there is no improving deviation in the neighbourhood of play. Bichler et al. are unable to prove that this entails a global equilibrium but provide empirical support for this stronger property across a suite of computational results. Second, while the expected payoff is a smooth function of the parameters of a strategy, a sample-averaging approach fails when used with backpropagation for gradient estimates. The problem is that for many auction designs of interest a bidder's payoff for a realization of values is a piecewise-constant function of its strategy and fails to yield useful gradients. To address this, Bichler et al. adopt an approach that perturbs the parameters of a policy and estimates the gradient empirically with respect to the change in payoff at these perturbed points. This yields the neural pseudogradient ascent (NPGA) algorithm. Simple ANNs with two hidden layers of ten units each are demonstrated to find a local equilibrium that also turns out to be a global equilibrium with high precision. New economic insights follow. For the LLG environment, for example, higher risk aversion on the part of the regional bidders is shown to improve the efficiency

and revenue of auctions by improving the coordination of bidders in the equilibrium.

The framework of Bichler et al. takes up the task of finding an equilibrium for a given auction design. A complementary approach seeks to design optimal auctions that are incentive compatible and have an equilibrium with a special structure. In particular, this looks for auctions with a truthful equilibrium in which bidders can do no better than bid their true values. ANNs have also been used to solve this design problem, where an ANN represents the allocation and pricing rules of an auction⁵. A question this sets up for future research is whether ANNs for automated auction design can be usefully combined with ANNs that can learn equilibria of auctions. Whereas the auctions in Bichler et al. are sealed-bid, with simultaneous reports, progress in this direction may expand the reach of automated design to other practically motivated designs such as those that proceed over multiple rounds of bidding and those that are robust to various kinds of collusive behaviour. □

David C. Parkes  

Harvard University, Cambridge, MA, USA.

 e-mail: parkes@eecs.harvard.edu

Published online: 9 August 2021

<https://doi.org/10.1038/s42256-021-00380-5>

References

1. Ausubel, L. M. & Baranov, O. *Int. J. Game Theory* **49**, 251–273 (2020).
2. Bichler, M., Fichtl, M., Heiderkrüger, S., Kohring, N. & Sutterer, P. *Nat. Mach. Intell.* <https://doi.org/10.1038/s42256-021-00365-4> (2021).
3. Jordan, J. S. *Games Econ. Behav.* **5**, 368–386 (1993).
4. Mazumdar, E., Ratliff, L. J. & Sastry, S. S. *SIAM J. Math. Data Sci.* **2**, 103–131 (2020).
5. Dütting, P., Feng, Z., Narasimhan, H., Parkes, D. C. & Ravindranath, S. S. In *Proc. 36th Int. Conf. Machine Learning* Vol. 97, 1706–1715 (PMLR, 2019).

Competing interests

The author declares no competing interests.