

STRATEGY SEARCHING IN STRATEGIC INTERACTION BETWEEN PARTIES

TECHNICAL FIELD

[0001] This specification relates to strategy searching in strategic interaction between two or more parties.

BACKGROUND

[0002] Strategic interaction between two or more parties can be modeled by a game that involves two or more parties (also referred to as players). In an Imperfect Information Game (IIG) that involves two or more players, a player only has partial access to the knowledge of her opponents before making a decision. This is similar to real-world scenarios, such as trading, traffic routing, and public auction. Many real life scenarios can be represented as IIGs, such as commercial competition between different companies, bidding relationships in auction scenarios, game relationships between a fraud party and an anti-fraud party.

[0003] Methods for solving an IIG is of great economic and societal benefits. Due to the hidden information, a player has to reason under the uncertainty about her opponents' information, and she also needs to act so as to take advantage of her opponents' uncertainty about her own information.

SUMMARY

[0004] Implementations of this specification include computer-implemented methods for strategy searching in strategic interaction between parties. More specifically, this specification describes examples of sampling schemes for performing a counterfactual regret minimization (CRF) algorithm in solving an imperfect information game (IIG), which can reduce the computational complexity and variance while improving the convergence speed of the CRF algorithm. This specification also describes techniques for performing counterfactual regret minimization (CRF) with neural networks, which can save memory space and provide faster convergence due to the generalization ability of the neural networks.

[0005] The subject matter described in this specification as implemented in particular embodiments realizes one or more of the following technical effects and advantages. In some embodiments, the described sampling techniques can help find better strategies of real-world scenarios such as resource allocation, product/service recommendation, cyber-attack prediction and/or prevention, traffic routing, fraud management, etc. that can be modeled or represented by strategic interaction between parties, such as, an IIG that involves two or more parties in a more efficient manner. In some embodiments, the described techniques can improve computational efficiency and reduce the computational load of counterfactual regret minimization (CRF) algorithm

in finding the best strategies of the real-world scenarios modeled by the IIG. In some embodiments, the described sampling techniques can provide a lower variance than outcome sampling, while being more memory efficient than external sampling. In some embodiments, the described techniques can improve the convergence speed of the CRF algorithm in finding Nash equilibrium for solving a game that represents one or more real-world scenarios. In some embodiments, the described techniques provide more balanced and comprehensive information of a game tree that represents the IIG so that the CRF algorithm can have a smaller variance and a faster convergence speed. In some embodiments, the described techniques save memory space and provide faster convergence by using neural networks in connection with the CRF algorithm. In some embodiments, the described techniques may only need a small amount of memory space for each iteration of the CRF algorithm.

[0006] This specification also provides one or more non-transitory computer-readable storage media coupled to one or more processors and having instructions stored thereon which, when executed by the one or more processors, cause the one or more processors to perform operations in accordance with embodiments of the methods provided herein.

[0007] This specification further provides a system for implementing the methods provided herein. The system includes one or more processors, and a computer-readable storage medium coupled to the one or more processors having instructions stored thereon which, when executed by the one or more processors, cause the one or more processors to perform operations in accordance with embodiments of the methods provided herein.

[0008] It is appreciated that methods in accordance with this specification may include any combination of the aspects and features described herein. That is, methods in accordance with this specification are not limited to the combinations of aspects and features specifically described herein, but also include any combination of the aspects and features provided.

[0009] The details of one or more embodiments of this specification are set forth in the accompanying drawings and the description below. Other features and advantages of this specification will be apparent from the description and drawings, and from the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] FIG. 1 is a diagram illustrating examples of partial game trees in one-card poker in accordance with embodiments of this specification.

[0011] FIG. 2 is a diagram illustrating examples of different sampling schemes in accordance with embodiments of this specification.

[0012] FIG. 3 is a pseudocode of an example of robust sampling Monte Carlo CFR (MCCFR) in accordance with embodiments of this specification.

[0013] FIG. 4 is a diagram illustrating an example of a double neural CFR algorithm applied to a game tree in accordance with embodiments of this specification.

[0014] FIG. 5 is a pseudocode of an example of a double neural CFR algorithm in accordance with embodiments of this specification.

[0015] FIG. 6 is a pseudocode of an example of an algorithm for optimizing a neural network in connection with a double neural CFR algorithm in accordance with embodiments of this specification.

[0016] FIG. 7 is a pseudocode of an example of a mini-batch MCCFR algorithm in accordance with embodiments of this specification.

[0017] FIG. 8 is a flow chart illustrating an example of a sampling process for performing MCCFR in accordance with embodiments of this specification.

[0018] FIG. 9 is a flow chart illustrating an example of a double neural CFR algorithm in accordance with embodiments of this specification.

[0019] FIG. 10 depicts a block diagram illustrating an example of a computer-implemented system used to provide computational functionalities associated with described algorithms, methods, functions, processes, flows, and procedures in accordance with embodiments of this specification.

[0020] FIG. 11 depicts examples of modules of an apparatus in accordance with embodiments of this specification.

[0021] FIG. 12 depicts examples of modules of another apparatus in accordance with embodiments of this specification.

[0022] Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0023] Implementations of this specification include computer-implemented methods for strategy searching in strategic interaction between parties, for example, by solving an imperfect information game (IIG). The IIG can represent one or more real-world scenarios such as resource allocation, product/service recommendation, cyber-attack prediction and/or prevention, traffic routing, fraud management, etc. that involves two or more parties (also referred to as players), where each party may have incomplete or imperfect information about the other party's decisions. More specifically, this specification describes examples of sampling schemes for performing a counterfactual regret minimization (CRF) algorithm in solving an IIG, which can reduce the computational complexity and variance while improving the convergence speed of the CRF algorithm. This specification also describes techniques for performing counterfactual regret minimization (CRF) with neural networks,

which can save memory space and provide faster convergence due to the generalization ability of the neural networks.

[0024] Nash equilibrium is a typical solution for an IIG that involves two or more players. Counterfactual Regret Minimization (CRF) is an algorithm designed to approximately find Nash equilibrium for large games. CRF tries to minimize overall counterfactual regret. It is proven that the average of the strategies in all iterations would converge to a Nash equilibrium. When solving a game, CRF in its original form (also referred to as original CRF, standard CRF, vanilla CFR, or simply, CRF) traverses the entire game tree in each iteration. Thus, the original CFR requires large memory for large, zero-sum extensive games such as heads-up no-limit Texas Hold'em. In some instances, the original CFR may not handle large games with limited memory.

[0025] A Monte Carlo CFR (MCCFR) was introduced to minimize counterfactual regret. The MCCFR can compute an unbiased estimation of counterfactual value and avoid traversing the entire game tree. Since only subsets of all information sets are visited in each iteration, MCCFR requires less memory than the original CFR.

[0026] MCCFR can be performed with an outcome sampling algorithm or an external sampling algorithm. The outcome sampling algorithm in MCCFR has a large variance, and it is difficult to converge to an approximate Nash equilibrium solution in fewer iteration steps. The external sampling algorithm in MCCFR has a smaller variance than the outcome sampling algorithm, but this method is similar to the disadvantage of CFR. When the game tree is large, it requires a very large memory space and cannot be extended to a complex large-scale IIG.

[0027] This specification discloses a robust sampling scheme. In the robust sampling scheme, each player uses a uniform sampling method to sample at a current decision point, and the other party samples according to a corresponding strategy. The reach probability corresponding to different iterations can be fixed. It can be proved that the robust sampling scheme has a smaller variance than the outcome sampling scheme in MCCFR, while being more memory efficient than the external sampling. In some embodiments, the robust sampling scheme can make the MCCFR solve Nash equilibrium with faster convergence.

[0028] This specification discloses a depth-dependent sampling scheme. The depth-dependent sampling scheme can allocate a higher sampling probability to a state closer to a terminal state than another state further from the terminal state (or closer to an initial or beginning state). In some embodiments, the depth-dependent sampling scheme can allow more states closer to a terminal state to be sampled, providing more well-around information of the IIG and thus improve the convergence rate of the MCCFR compared to existing sampling schemes.

[0029] This specification further discloses a double neural CFR algorithm. Existing CFR methods such as CFR and MCCFR use two large tabular-based memories to record the cumulative regret and average strategy for all information sets. Such tabular representation makes these methods difficult to apply to large extensive-form games with limited time and space.

[0030] By contrast, the double neural CFR algorithm uses two neural networks to compute approximate Nash equilibrium of IIG. For example, one of the neural networks can be used to learn the cumulative regret and the other one can be used to learn the cumulative numerator of the average strategy profile. With the help of these two networks, the double neural CFR algorithm does not need to use two large tabular-based memories. Based on the generalization ability of the compact neural network, the cumulative regret and the average strategy can be learned and produced. The disclosed double neural CFR algorithm can keep the benefit of MCCFR in requiring less computational load yet without the need for two large tabular memories. The disclosed double neural CFR algorithm can be used in large games even with memory constraints. In some embodiments, the double neural method can achieve a lower exploitability with fewer iterations than existing techniques. In addition, in some embodiments, the double neural CFR can also continually improve after initialization from a poor tabular strategy.

[0031] In some embodiments, the described techniques can be used, for example, in AI poker, recommendation platforms, and many other AI and machine learning applications. The described techniques use the Monte Carlo method and does not require variables for the entire game tree.

[0032] In some embodiments, an extensive-form game with a finite set $N = \{0, 1, \dots, n-1\}$ of players can be represented as follows. Define h_i^v as a hidden variable of player i in an IIG. For example, in a poker game, h_i^v can refer to the private cards of player i . H refers to a finite set of histories. Each member $h = (h_i^v)_{i=0,1,\dots,n-1}(a_l)_{l=0,\dots,L-1} = h_0^v h_1^v \dots h_{n-1}^v a_0 a_1 \dots a_{L-1}$ of H denotes a possible history (or state), which includes each player's hidden variable and L actions taken by players including chance. For player i , h also can be denoted as $h_i^v h_{-i}^v a_0 a_1 \dots a_{L-1}$, where h_{-i}^v refers to the opponent's hidden variables. The empty sequence $\emptyset$ is a member of H . The expression $h_j \sqsubseteq h$ denotes that h_j is a prefix of h , where $h_j = (h_i^v)_{i=0,1,\dots,n-1}(a_l)_{l=1,\dots,L'-1}$ and $0 < L' < L$. $Z \subseteq H$ denotes the terminal histories and any member $z \in Z$ is not a prefix of any other sequences. $A(h) = \{a : ha \in H\}$ is the set of available actions after non-terminal history $h \in H \setminus Z$. A player function P assigns a member of $N \cup \{c\}$ to each non-terminal history, where c denotes the chance player identifier (ID), which typically can be, for example, -1. $P(h)$ is the player who takes an action after history h .

[0033] I_i of a history $\{h \in H : P(h) = i\}$ is an information partition of player i . A set $I_i \in \mathcal{I}_i$ is an information set of player i . $I_i(h)$ refers to information set I_i at state h . In some embodiments, I_i could

only remember the information observed by player i including player i 's hidden variable and public actions. Therefore I_i indicates a sequence in the IIG, i.e., $h_i^v a_0 a_2 \dots a_{L-1}$. In some embodiments, for $I_i \in \mathcal{I}_i$ and for any $h \in I_i$, the set $A(h)$ can be denoted by $A(I_i)$ and the player $P(h)$ is denoted by $P(I_i)$. For each player $i \in N$, a utility function $u_i(z)$ defines a payoff of a terminal state z . A more detailed explanation of these notations and definitions will be discussed below including an example shown in FIG. 1.

[0034] FIG. 1 is a diagram 100 illustrating examples of partial game trees 102 and 104 in One-Card Poker in accordance with embodiments of this specification. One-Card Poker is a two-players IIG of poker. One-Card Poker is an example of an extensive-form game. The game rules are defined as follows. Each player is dealt one card from a deck of X cards. The first player can pass or bet. If the first player bet, the second player can call or fold. If the first player pass, the second player can pass or bet. If second player bet, the first player can fold or call. The game ends with two passes, a call, or a fold. The fold player will lose 1 chips. If the game ended with two passes, the player with higher card win 1 chips. If the game end with a call, the player with higher card win 2 chips.

[0035] A game tree is a directed graph. The nodes of the game tree represent positions (or states of a player) in a game and of the game tree represent can represent moves or actions of a player of the game. In FIG. 1, z_i denotes a terminal node, representing a terminal state, and h_i denotes a non-terminal node. Each of the partial game trees 102 and 104 has a root node h_0 representing a chance. There are 19 distinct nodes in the first partial tree 102, corresponding to 9 non-terminal nodes h_i including chance h_0 and 10 terminal nodes z_i in the left tree.

[0036] In the first partial tree 102, two players (player 0 and player 1) are dealt (queen, jack) as shown as "0:Q 1:J" in the left subtree and (queen, king) as shown as "0:Q 1:K" in the right subtree.

[0037] The trajectory from the root node to each node is a history of actions. Actions are presented by letters (e.g., F, C, P, and B) or representations (e.g., "0:Q 1:J") next to edges (denoted by arrows) of the game tree. The letters F, C, P, B refer to fold, call, pass, and bet, respectively.

[0038] In an extensive-form game, h_i refers to the history of actions. For example, as illustrated in the first partial tree 102, h_3 includes actions 0:Q, 1:J and P. h_7 includes actions 0:Q, 1:J, P and B. h_8 includes actions 0:Q, 1:K, P and B. In the first partial tree 102, $h_3 \sqsubseteq h_7$, that is, h_3 is a prefix of h_7 . $A(h_3) = \{P, B\}$ indicating that the set of available actions after non-terminal history h_7 are P and B. $P(h_3) = 1$ indicating that the player who takes an action after history h_3 is player 1.

[0039] In the IIG, the private card of player 1 is invisible to player 0, therefore h_7 and h_8 are actually the same for player 0. An information set can be used to denote the set of these

undistinguished states. Similarly, h_1 and h_2 are in the same information set. For the right partial tree 104, h'_3 and h'_5 are in the same information set; h'_4 and h'_6 are in the same information set.

[0040] Typically, any $I_i \in \mathcal{I}$ could only remember the information observed by player i including player i 's hidden variables and public actions. For example, as illustrated in the first partial tree 102, the information set of h_7 and h_8 indicates a sequence of 0:Q, P, and B. Because h_7 and h_8 are undistinguished by player 0 in the IIG, if I_0 is the information set of h_7 and h_8 , $I_0 = I_0(h_7) = I_0(h_8)$.

[0041] A strategy profile $\sigma = \{\sigma_i | \sigma_i \in \Sigma_i, i \in N\}$ is a collection of strategies for all players, where Σ_i is the set of all possible strategies for player i . σ_{-i} refers to strategy of all players other than player i . For player $i \in N$, the strategy $\sigma_i(I_i)$ is a function, which assigns an action distribution over $A(I_i)$ to information set I_i . $\sigma_i(a|h)$ denotes the probability of action a taken by player $i \in N \cup \{c\}$ at state h . In an IIG, if two or more states have the same information set, the two or more states have a same strategy. That is, $\forall h_1, h_2 \in I_i, I_i = I_i(h_1) = I_i(h_2), \sigma_i(I_i) = \sigma_i(h_1) = \sigma_i(h_2), \sigma_i(a|I_i) = \sigma_i(a|h_1) = \sigma_i(a|h_2)$. For example, I_0 is the information set of h_7 and h_8 , $I_0 = I_0(h_7) = I_0(h_8), \sigma_0(I_0) = \sigma_0(h_7) = \sigma_0(h_8), \sigma_0(a|I_0) = \sigma_0(a|h_7) = \sigma_0(a|h_8)$. In FIG. 1, the same color other than gray for each state in the same information set.

[0042] For player i , the expected game utility of the strategy profile σ is denoted as $u_i^\sigma = \sum_{z \in Z} \pi^\sigma(z) u_i(z)$, which is the expected payoff of all possible terminal nodes. Given a fixed strategy profile σ_{-i} , any strategy $\sigma_i^* = \arg \max_{\sigma'_i \in \Sigma_i} u_i^{(\sigma'_i, \sigma_{-i})}$ of player i that achieves maximize payoff against π_{-i}^σ is a best response. For two players' extensive-form games, a Nash equilibrium is a strategy profile $\sigma^* = (\sigma_0^*, \sigma_1^*)$ such that each player's strategy is a best response to the opponent. An ϵ -Nash equilibrium is an approximation of a Nash equilibrium, whose strategy profile σ^* satisfies: $\forall i \in N, u_i^{\sigma_i^*} + \epsilon \geq \max_{\sigma'_i \in \Sigma_i} u_i^{(\sigma'_i, \sigma_{-i}^*)}$.

[0043] Exploitability of a strategy σ_i can be defined as $\epsilon_i(\sigma_i) = u_i^{\sigma_i^*} - u_i^{(\sigma_i, \sigma_{-i}^*)}$. A strategy is unexploitable if $\epsilon_i(\sigma_i) = 0$. In large two player zero-sum games such as poker, $u_i^{\sigma_i^*}$ can be intractable to compute. However, if the players alternate their positions, the value of a pair of games is zero, i.e., $u_0^{\sigma_0^*} + u_1^{\sigma_1^*} = 0$. The exploitability of strategy profile σ can be defined as $\epsilon(\sigma) = (u_1^{(\sigma_0, \sigma_1^*)} + u_0^{(\sigma_0^*, \sigma_1)})/2$.

[0044] For iterative methods such as CFR, σ^t can refer to the strategy profile at the t -th iteration. The state reach probability of history h can be denoted by $\pi^\sigma(h)$ if players take actions according to σ . For an empty sequence $\pi^\sigma(\emptyset) = 1$. The reach probability can be decomposed into $\pi^\sigma(h) = \prod_{i \in N \cup \{c\}} \pi_i^\sigma(h) = \pi_i^\sigma(h) \pi_{-i}^\sigma(h)$ according to each player's contribution, where

$$\pi_i^\sigma(h) = \prod_{h' \sqsubseteq h, P(h')=P(h)} \sigma_i(a|h') \text{ and } \pi_{-i}^\sigma(h) = \prod_{h' \sqsubseteq h, P(h') \neq P(h)} \sigma_{-i}(a|h').$$

[0045] The reach probability of information set I_i (also referred to as information set reach probability) can be defined as $\pi^\sigma(I_i) = \sum_{h \in I_i} \pi^\sigma(h)$. If $h' \sqsubseteq h$, the interval state reach probability from state h' to h can be defined as $\pi^\sigma(h', h)$, then $\pi^\sigma(h', h) = \pi^\sigma(h)/\pi^\sigma(h')$. The reach probabilities $\pi_i^\sigma(I_i), \pi_{-i}^\sigma(I_i), \pi_i^\sigma(h', h)$, and $\pi_{-i}^\sigma(h', h)$ can be defined similarly.

[0046] In large and zero-sum IIGs, CFR is proved to be an efficient method to compute Nash equilibrium. It is proved that the state reach probability of one player is proportional to posterior probability of the opponent's hidden variable, *i.e.*, $p(h_{-i}^v | I_i) \propto \pi_{-i}^\sigma(h)$, where h_{-i}^v and I_i indicate a particular h .

[0047] For player i and strategy profile σ , the counterfactual value (CFV) $v_i^\sigma(h)$ at state h can be define as

$$v_i^\sigma(h) = \sum_{h \sqsubseteq z, z \in Z} \pi_{-i}^\sigma(h) \pi^\sigma(h, z) u_i(z) = \sum_{h \sqsubseteq z, z \in Z} \pi_i^\sigma(h, z) u'_i(z) \quad (1)$$

where $u'_i(z) = \pi_{-i}^\sigma(z) u_i(z)$ is the expected reward of player i with respect to the approximated posterior distribution of the opponent's hidden variable. The action counterfactual value of taking action a can be denoted as $v_i^\sigma(alh) = v_i^\sigma(ha)$ and the regret of taking this action is $r_i^\sigma(alh) = v_i^\sigma(alh) - v_i^\sigma(h)$.

[0048] Similarly, the CFV of information set I_i can be defined as $v_i^\sigma(I_i) = \sum_{h \in I_i} v_i^\sigma(h)$ and the regret of action a given the information set I_i can be defined as $r_i^\sigma(a|I_i) = \sum_{z \in Z, ha \sqsubseteq z, h \in I_i} \pi_i^\sigma(ha, z) u'_i(z) - \sum_{z \in Z, h \sqsubseteq z, h \in I_i} \pi_i^\sigma(h, z) u'_i(z)$ (1a)

[0049] Then the cumulative regret of action a after T iterations can be calculated according to Eq. (2):

$$R_i^T(a|I_i) = \sum_{t=1}^T \left(v_i^{\sigma^t}(a|I_i) - v_i^{\sigma^t}(I_i) \right) = R_i^{T-1}(a|I_i) + r_i^{\sigma^T}(a|I_i) \quad (2)$$

where $R_i^0(a|I_i) = 0$. Define $R_i^{T,+}(a|I_i) = \max(R_i^T(a|I_i), 0)$, the current strategy (or behavior strategy) at $T + 1$ iteration can be updated, for example, based on regret matching, according to Eq. (3) below:

$$\sigma_i^{T+1}(a|I_i) = \begin{cases} \frac{R_i^{T,+}(a|I_i)}{\sum_{a \in A(I_i)} R_i^{T,+}(a|I_i)}, & \text{if } \sum_{a \in A(I_i)} R_i^{T,+}(a|I_i) > 0 \\ \frac{1}{|A(I_i)|}, & \text{otherwise} \end{cases} \quad (3).$$

[0050] The average strategy $\bar{\sigma}_i^T$ from iteration 1 to T can be defined as:

$$\bar{\sigma}_i^T(a|I_i) = \frac{\sum_{t=1}^T \pi_i^{\sigma^t}(I_i) \sigma_i^t(a|I_i)}{\sum_{t=1}^T \pi_i^{\sigma^t}(I_i)} \quad (4)$$

where $\pi_i^{\sigma^t}(I_i)$ denotes the information set reach probability of I_i at t -th iteration and is used to weigh the corresponding current strategy $\sigma_i^t(a|I_i)$.

[0051] Define $s^t(a|I_i) = \pi_i^{\sigma^t}(I_i)\sigma_i^t(a|I_i)$ as an additional numerator in iteration t , then the cumulative numerator of the average strategy $\bar{\sigma}_i^T$ can be defined as

$$S^T(a|I_i) = \sum_{t=1}^T \pi_i^{\sigma^t}(I_i)\sigma_i^t(a|I_i) = S^{T-1}(a|I_i) + s_i^T(a|I_i), (5)$$

where $S^0(a|I_i) = 0$.

[0052] When solving a game, the original CFR traverses the entire game tree in each iteration. Thus, the original CFR may not handle large games with limited memory. A Monte Carlo CFR (MCCFR) was introduced to minimize counterfactual regret. The MCCFR can compute an unbiased estimation of counterfactual value and avoid traversing the entire game tree. Since only subsets of all information sets are visited in each iteration, MCCFR requires less memory than the original CFR.

[0053] For example, define $Q = \{Q_1, Q_2, \dots, Q_m\}$, where $Q_j \in Z$ is a block of sampling terminal histories in each iteration, such that Q_j spans the set Z . Generally, different Q_j may have an overlap according to a specified sampling scheme. Several sampling schemes can be used.

[0054] FIG. 2 is a diagram 200 illustrating examples of different sampling schemes in accordance with embodiments of this specification. Specifically, the subplot A illustrates an example of an external sampling scheme 202 of a game tree; the subplot B illustrates an example of an outcome sampling scheme 204 of a game tree, and the subplot C illustrates an example of a robust sampling scheme 206 of a game tree.

[0055] As illustrated in FIG. 2, a circle represents a player 0 node, a rectangular represents a player 1 node, and a triangle represents a chance node. The solid edges or arrows present sampled actions whereas the dashed edges or arrows present non-sampled actions. The shaded nodes present sampled nodes whereas the blank nodes present non-sampled nodes.

[0056] Take the update of player 0 as an example, with the external sampling scheme 202 as shown in subplot A, the player 0 node traverses all branches of the player 0 node, a non-player 0 node (e.g., player 1 node and chance node) randomly samples a branch according to a corresponding sampling strategy.

[0057] The outcome sampling scheme does not distinguish different players. As shown in subplot B, the outcome sampling scheme 204 randomly samples one branch for all players according to the corresponding sampling strategy. As such, only one trajectory will be sampled under the outcome sampling scheme.

[0058] As shown in subplot C, the robust sampling scheme 206 randomly selects k branches according to a uniform distribution for player 0, and performs random sampling on one branch for a non-player 0 node according to a corresponding sampling strategy. By varying the value of k , the robust sampling scheme can sample multiple paths, or a single path, for example, depending on

actual memory needs or system specification. Unlike the external sampling scheme, the robust sampling scheme does not require the knowledge of all possible actions and the variables at the current player i 's decision point each time.

[0059] In some embodiments, in the external sampling and outcome sampling schemes, each block $Q_j \in Q$ represents a partition of Z . Define q_{Q_j} as the probability of considering block Q_j , where $\sum_{j=1}^m q_{Q_j} = 1$. Define $q(z) = \sum_{j: z \in Q_j} q_{Q_j}$ as the probability of considering a particular terminal history z . In some embodiments, vanilla CFR can be considered as a special case of MCCFR, where $Q = \{Z\}$ only contain one block and $q_{Q_1} = 1$.

[0060] In the outcome sampling scheme, only one trajectory will be sampled, such that $\forall Q_j \in Q$, $|Q_j| = 1$ and $|Q| = |Z|$. For information set I_i , a sample estimate of counterfactual value is $\tilde{v}_i^\sigma(I_i|Q_j) = \sum_{h \in I_i, z \in Q_j, h \sqsubseteq z} \frac{1}{q(z)} \pi_{-i}^\sigma(z) \pi_i^\sigma(h, z) u_i(z)$.

[0061] It is proved that the sampling counterfactual value in MCCFR is the unbiased estimation of actual counterfactual value in CFR: $E_{j \sim q_{Q_j}} [\tilde{v}_i^\sigma(I_i|Q_j)] = v_i^\sigma(I_i)$.

[0062] Define σ^{rs} as a sampling strategy profile, where σ_i^{rs} is the sampling strategy for player i and σ_{-i}^{rs} are the sampling strategies for players except player i . In some embodiments, for both external sampling and outcome sampling, $\sigma_{-i}^{rs} = \sigma_{-i}$. The regret of the sampled action $a \in A(I_i)$ can be defined as:

$$\tilde{r}_i^\sigma((a|I_i)|Q_j) = \sum_{z \in Q_j, ha \sqsubseteq z, h \in I_i} \pi_i^\sigma(ha, z) u_i^{rs}(z) - \sum_{z \in Q_j, h \sqsubseteq z, h \in I_i} \pi_i^\sigma(h, z) u_i^{rs}(z), \quad (6)$$

where $u_i^{rs}(z) = \frac{u_i(z)}{\pi_i^{\sigma^{rs}}(z)}$ is a new utility weighted by $\frac{1}{\pi_i^{\sigma^{rs}}(z)}$. The sample estimate for cumulative regret of action a after T iterations can be defined as

$$\tilde{R}_i^T((a|I_i)|Q_j) = \tilde{R}_i^{T-1}((a|I_i)|Q_j) + \tilde{r}_i^{\sigma^T}((a|I_i)|Q_j) \text{ with } \tilde{R}_i^0((a|I_i)|Q_j) = 0.$$

[0063] For robust sampling, the sampling profile can be defined as $\sigma^{rs(k)} = (\sigma_i^{rs(k)}, \sigma_{-i})$, where player i can randomly select k actions according to sampling strategy $\sigma_i^{rs(k)}(I_i)$ for each information set I_i and other players can randomly select one action according to strategy σ_{-i} .

[0064] In some embodiments, if player i randomly selects $\min(k, |A(I_i)|)$ actions according to discrete uniform distribution $\text{unif}(0, |A(I_i)|)$ at information set I_i . That is, $\sigma_i^{rs(k)}(a|I_i) = \frac{\min(k, |A(I_i)|)}{|A(I_i)|}$, then the reach probability of information set I_i if player i takes actions according to sampling strategy or profile $\sigma_i^{rs(k)}$ can be calculated by:

$$\pi_i^{\sigma^{rs(k)}}(I_i) = \prod_{h \in I_i, h' \sqsubseteq h, h' a \sqsubseteq h, h' \in I_i'} \frac{\min(k, |A(I_i')|)}{|A(I_i')|},$$

and the weighted utility $u_i^{rs(k)}(z)$ can be a constant number in each iteration given the sampling profile $\sigma^{rs(k)}$, which has a low variance. In addition, because the weighted utility no longer requires explicit knowledge of the opponent's strategy, robust sampling can be used for online regret minimization.

[0065] To simplify notations, let $k = \max$ refer to $k = \max_{I_i \in \mathcal{I}} |A(I_i)|$. If $k = \max$ and player i randomly selects k actions according to a discrete uniform distribution $\text{unif}(0, |A(I_i)|)$ at information set I_i , $\forall I_i \in \mathcal{I}, \forall I_i \in \mathcal{I}, \forall a \in A(I_i), \sigma_i^{rs(k)}(a|I_i) \sim \text{unif}(0, |A(I_i)|)$, then robust sampling can be similar to external sampling when $k = \max_{I_i \in \mathcal{I}} |A(I_i)|$.

[0066] If $k = 1$ and $\sigma_i^{rs(k)} = \sigma_i$, only one history z is sampled in this case, then $u_i^{rs(1)}(z) = \frac{u_i(z)}{\pi_{\sigma^{rs(k)}}(z)}$, $\exists h \in I_i$, for $a \in A^{rs(k)}(I_i)$,

$$\begin{aligned} \tilde{r}_i^\sigma((a|I_i)|Q_j) &= \tilde{r}_i^\sigma((a|h)|Q_j) = \sum_{z \in Q_j, ha \sqsubseteq z, h \in I_i} \pi_i^\sigma(ha, z) u_i^{rs}(z) - \sum_{z \in Q_j, h \sqsubseteq z, h \in I_i} \pi_i^\sigma(h, z) u_i^{rs}(z) \\ &= \frac{(1 - \sigma_i(a|h)) u_i(z)}{\pi_i^\sigma(ha)} \end{aligned}$$

[0067] If action a is not sampled at state h , i.e., $a \notin A^{rs(k)}(I_i)$, the regret is $\tilde{r}_i^\sigma((a|h)|j) = 0 - \tilde{v}_i^\sigma(h|j)$. In this case, robust sampling is similar to the outcome sampling when $k = 1$ and $\sigma_i^{rs(k)} = \sigma_i$.

[0068] If $k = 1$, and player i randomly selects one action according to discrete uniform distribution $\text{unif}(0, |A(I_i)|)$ at information set I_i . Then robust sampling can be similar to outcome sampling. For example, if $k = 1$, and player i randomly selects one action according to a discrete uniform distribution $\text{unif}(0, |A(I_i)|)$ at information set I_i , then $u_i^{rs(1)}(z) = \frac{u_i(z)}{\pi_{\sigma^{rs(k)}}(z)}$ is a constant, $\exists h \in I_i$, for $a \in A^{rs(k)}(I_i)$,

$$\begin{aligned} \tilde{r}_i^\sigma((a|I_i)|Q_j) &= \sum_{z \in Q_j, ha \sqsubseteq z, h \in I_i} \pi_i^\sigma(ha, z) u_i^{rs}(z) - \sum_{z \in Q_j, h \sqsubseteq z, h \in I_i} \pi_i^\sigma(h, z) u_i^{rs}(z) \\ &= (1 - \sigma_i(a|h)) \pi_i^\sigma(ha, z) u_i^{rs(1)}(z). \end{aligned}$$

[0069] If action a is not sampled at state h , i.e., $a \notin A^{rs(k)}(I_i)$, the regret is $\tilde{r}_i^\sigma((a|h)|j) = 0 - \tilde{v}_i^\sigma(h|j)$. Compared to outcome sampling, the robust sampling in this case has a lower variance because of the constant $u_i^{rs(1)}(z)$.

[0070] FIG. 3 is a pseudocode 300 of an example of robust sampling MCCFR in accordance with embodiments of this specification. As shown in lines 1-5 of the pseudocode 300, the overall robust sampling MCCFR is an iterative algorithm with an input of the total number of iterations, t . Within

each iteration t , a function robust sampling MCCFR (RS-MCCFR) is called for player 0 and player 1 (as shown in lines 3 and 4) to update the cumulative regret $R^t(\cdot | I_i)$ and average strategy numerator $S^t(a | I_i)$. The function RS-MCCFR can be defined as shown in lines 6-30 of the pseudocode 300. The function RS-MCCFR return the counterfactual value of each information set as the output. In some embodiments, the counterfactual value can be used to calculate the counterfactual regret. As such, the cumulative regret and time-average strategy can be obtained accordingly.

[0071] Specifically, the function RS-MCCFR samples the actions according to the robust sampling scheme as described above in connection with FIG. 2. As shown in line 16 of the pseudocode 300, k different actions can be sampled and collected as $A^{rs(k)}(I_i)$ according to the robust sampling strategy $\sigma_i^{rs(k)}$.

[0072] In some embodiments, a depth-dependent sampling scheme can be used to provide more balanced or well-around information of a game represented by a game tree. For example, the sampling strategy $\sigma_i^{rs}(a|h)$ can be a function of a depth of the state h in a game tree. For example, the sampling strategy $\sigma_i^{rs}(a|h)$ can be designed such that a state closer to the terminal state has a higher probability of being sampled than a state closer to the initial state (e.g., represented by the root of the game tree). As an example, the depth-dependent sampling scheme can be implemented by applying different weights to sampling probabilities of different states with different depths. Such a depth-dependent sampling can help provide more information on states closer to the terminal states, which can be advantageous because typically there are more states closer to the terminal states than those closer to the initial state (e.g., due to the branching structure of the game tree) and these nodes have a lower chance of being sampled than nodes closer to the root node in sampled trajectories under existing sampling schemes.

[0073] In some embodiments, the depth-dependent sampling scheme can be used in combination with the robust sampling, outcome sampling, external sampling, or any other suitable sampling algorithms. For example, the depth-dependent sampling scheme can further improve the variance and convergence speed of either of the robust sampling, outcome sampling, and external sampling as the latter three sampling schemes focus more on horizontal sampling among different actions of a state of a player (e.g., represented by the different branches of a node of a game tree).

[0074] FIG. 4 is a diagram illustrating an example 400 of a double neural CFR algorithm applied to a game tree 410 in accordance with embodiments of this specification. The double neural CFR algorithm 400 use two neural networks 420 and 430 to compute approximate Nash equilibrium of an IIG such as represented by the game tree 410. As illustrated in FIG. 4, one neural network 420 is used for obtaining a cumulative regret and is referred to as RegretSumNetwork (RSN). The other

neural networks 430 is used for obtaining an average strategy and is referred to as AveStrategyNetwork (ASN).

[0075] In some embodiments, the iterative updates of the CFR algorithm maintain two strategies: the current strategy $\sigma_i^t(a|I_i)$, and the average strategy $\bar{\sigma}_i^t(a|I_i)$ for $\forall i \in N, \forall I_i \in I_i, \forall a \in A(I_i), \forall t \in \{1, \dots, T\}$. Accordingly, the two neural networks 420 and 430 can be designed to record these two strategies, respectively, in an iterative fashion. In some embodiments, the example 400 of the double neural CFR algorithm can be referred to a double incremental CFR algorithm as the neural networks are trained or optimized based on new, additional samples in each iteration.

[0076] According to Eq. (3), a current strategy $\sigma^{t+1}(a|I_i)$ can be computed by the cumulative regret $R^t(a|I_i)$. In some embodiments, only the numerator in Eq. (3) is tracked since the normalization in the denominator can easily be computed when the strategy is used. Given information set I_i and action a , the neural network RSN 420, denoted as $\mathcal{R}(a, I_i | \theta_R^t)$, can be used to learn $R^t(a|I_i)$, where θ_R^t is the parameter in the RSN 420 at t -th iteration.

[0077] As shown FIG. 4, a memory M_R^t can be defined as $M_R^t = \{(I_i, r_i^{\sigma^t}((a|I_i)|Q_j)) | \forall i \in N, \forall a \in A(I_i), h \in I_i, h \sqsubseteq z, z \in Q_j\}$. Each member of M_S^t can include visited information set I_i and the corresponding regret $\tilde{r}_i^{\sigma^t}((a|I_i)|Q_j)$, where Q_j is the sampled block in t -th iteration. According to Eq. (2), $\mathcal{R}(a, I_i | \theta_R^{t+1})$ can be estimated using the following optimization:

$$\theta_R^{t+1} \leftarrow \operatorname{argmin}_{\theta_R^{t+1}} \sum_{(I_i, \tilde{r}_i^{\sigma^t}((a|I_i)|Q_j)) \in M_R^t} (\mathcal{R}(a, I_i | \theta_R^t) + \tilde{r}_i^{\sigma^t}((a|I_i)|Q_j) - \mathcal{R}(a, I_i | \theta_R^{t+1}))^2, \quad (7)$$

According to Eq. (4), the approximate Nash equilibrium is the weighted average of all previous strategies over T iterations. Similar to the cumulative regret, another neural network ASN 430, denoted as $S(a, I_i | \theta_S^t)$, can be used to learn the numerator of the average strategy. Define another memory M_S^t as

$$M_S^t = \{(I_i, s^t(a|I_i)) | \forall i \in N, \forall a \in A(I_i), h \in I_i, h \sqsubseteq z, z \in Q_j\} = \{(I_i, \pi_i^{\sigma^t}(I_i) \sigma_i^t(a|I_i)) | \forall i \in N, \forall a \in A(I_i), h \in I_i, h \sqsubseteq z, z \in Q_j\}.$$

Each member of M_S^t can include the visited information set I_i and the value of $\pi_i^{\sigma^t}(I_i) \sigma_i^t(a|I_i)$, where Q_j is the sampled block in t -th iteration. Then the parameter θ_S^{t+1} of the ASN 430 can be estimated by the following formula:

$$\theta_S^{t+1} \leftarrow \operatorname{argmin}_{\theta_S^{t+1}} \sum_{(I_i, s_i^t(a|I_i)) \in M_S^t} (S(a, I_i | \theta_S^t) + s_i^t(a|I_i) - S(a, I_i | \theta_S^{t+1}))^2, \quad (8)$$

[0078] In some embodiments, in each iteration, both $\mathcal{R}(a, I_i | \theta_R^t)$ and $S(a, I_i | \theta_S^t)$ can be optimized by optimizing the objectives in Eq. (7) and Eq. (8) according to a gradient descent method, such as a mini-batch stochastic gradient descent method described with respect to FIG. 7 below.

[0079] In some embodiments, the average strategy does not need to update in each iteration if a large memory is available to aggregate and save $s_i^t(a|I_i)$ within multiple iterations. If the memory M_S^t is full, the incremental value can be learned by optimizing Eq. (8).

[0080] In some embodiments, in each iteration, only a small subset of information sets are sampled, which may lead to that the neural network RSN 420 and ASN 430 forget values of those unobserved or un-sampled information sets. To address this issue, the neural network parameters from a previous iteration can be used as the initialization of a current iteration, which gives an online learning/adaptation flavor to the updates. Furthermore, due to the generalization ability of the neural networks, even samples from a small number of information sets can be used to update the new neural networks, and the newly updated neural networks can produce good values for the cumulative regret and the average strategy In some embodiments.

[0081] In some embodiments, as the number of iterations t increases, the value of $R_i^t(a|I_i)$ may become increasingly large, potentially making it difficult for the neural network to learn. To address this issue, the cumulative regret can be normalized by a factor of $\sqrt{t}$ to make its range more stable. This can be understood from the regret bound of online learning. More specifically, let $\Delta = \max_{I_i, a, t} |R_i^t(a|I_i) - R_i^{t-1}(a|I_i)|, \forall I_i \in \mathcal{I}, a \in A(I_i), t \in \{1, \dots, T\}$. $R_i^t(a|I_i) \leq \Delta \sqrt{|A|t}$, where $|A| = \max_{I_i \in \mathcal{I}} |A(I_i)|$. In some embodiments, the neural network RSN 420 can be used to track $\hat{R}_i^t(a|I_i) = R_i^t(a|I_i)/\sqrt{t}$, and update it by

$$\hat{R}_i^t(a|I_i) = \frac{\sqrt{t-1}\hat{R}_i^{t-1}(a|I_i)}{\sqrt{t}} + \frac{r_i^{\sigma^t}(a|I_i)}{\sqrt{t}}, \quad (9)$$

where $\hat{R}_i^0(a|I_i) = 0$.

[0082] In some embodiments, in the double incremental CFR algorithm, the memories M_R^t and M_S^t may be cleared after each iteration, for example, due to the limited size of the memories M_R^t and M_S^t . In some embodiments, for example in a large game, even with the double incremental CFR algorithm that uses the neural networks to learn the cumulative regrets and the average strategy, the size of the memories M_R^t and M_S^t may still need to be very large to record the cumulative regret and average strategies for each iteration.

[0083] In some embodiments, to continually improve the average strategy with limited memory but infinity iterations and/or to further relieve the requirement on the memory size, a double reservoir CFR algorithm can be used that use two reservoirs M_R and M_S to save the sampled cumulative regrets and the average strategies across different iterations and to dynamically learn the cumulative regret and average strategy.

[0084] In some embodiments, an average cumulative regret after T iterations can be obtained according to Eq. (10) by rewriting Eq. (2):

$$\bar{R}_i^T(a|I_i) = R_i^T(a|I_i)/T \quad (10)$$

[0085] Similarly, the average strategy can be the normalization of cumulative strategy as shown in Eq. (4), which is weighted average strategy $\sigma_i^t(a|I_i)$ by its reach probability $\pi_i^{\sigma^t}(I_i)$.

[0086] In some embodiments, two uniform reservoirs M_R and M_S can be used to save the sampled $\tilde{r}_i^{\sigma^t}((a|I_i)|Q_j)$ and $s^t(a|I_i)$, respectively. Specifically, M_R can be the reservoir to save samples in M_R^t and M_S can be the reservoir to save sampled in M_S^t . New samples can be inserted into the reservoir by a reservoir sampling algorithm. Reservoir sampling includes a family of randomized algorithms for randomly choosing k items from a list containing n items. For example, if the reservoir is not full, new samples can be added to the reservoir directly. If a reservoir is full, new samples can replace old samples according to, for example, a first in first out (FIFO) principle or according to a uniform random distribution or another distribution.

[0087] Note that both the double incremental CFR algorithm and the double reservoir CFR algorithm employ the ideas in online learning and use two neural networks to learn the updating regret and average strategy, respectively. In some embodiments, the ASN does not need to be updated in each iteration while the RSN may need to be optimized after Monte Carlo sampling so as to produce a new behavior strategy. For example, when a new behavior strategy is used to traverse the game tree, the RSN may need to be updated each iteration. On the other hand, the ASN can be used as the final approximated Nash equilibrium which is the weighted average of behavior strategy. The ASN can serve as the output of the double neural CFR algorithm. If there is a large enough data store to save all the samples, it is only needed to optimize the average strategy at the last step. In practice, for a large game, the large data store can be very expensive. As such, the average strategy can be incrementally optimized if the data store (e.g., the reservoir M_S^t) is full. As such, the double neural CFR algorithm can include two variations. In the double incremental algorithm, the neural network (e.g., one or both RSN and ASN) is optimized by the incremental samples only, while the double reservoir algorithms, the neural network (e.g., one or both RSN and ASN) can be optimized by all the samples in the reservoirs.

[0088] The double neural CFR algorithm and the double reservoir CFR algorithm have different sample collections. For the double incremental CFR, the neural network is optimized based on the newly added samples. For double reservoir CFR, the neural network is optimized based on the samples in the fixed-size reservoirs. In addition, in the double reservoir method, the average strategy can be optimized by maximum log-likelihood rather than minimum squared error.

[0089] FIG. 5 is a pseudocode 500 of an example of a double neural CFR algorithm in accordance with embodiments of this specification. The example of the double neural CFR algorithm includes an option of using the double neural CFR algorithm or the double reservoir CFR algorithm.

[0090] Lines 3-7 of the pseudocode 500 show examples of initialization strategies in the first iteration. For example, if the system warm starts from an existing CFR method (e.g., tabular-based CFR or MCCFR methods or the double neural CFR method), the neural networks can be initialized from an existing strategy profile to clone the cumulative regrets and strategy. If there is no warm start initialization, the double neural CFR algorithm can start by randomly initializing the parameters in RSN and ASN at iteration $t = 1$.

[0091] In some embodiments, if the double incremental CFR algorithm is used, as shown in line 8 of the pseudocode 500, sampling methods can be used return the counterfactual regret and the numerator of average strategy for the sampled information sets in this iteration. The counterfactual regret and the numerator of average strategy for the sampled information sets in this iteration can be saved in memories M_R^t and M_S^t respectively. The sampling methods can include, for example, a mini-batch robust sampling method algorithm described with respect to FIG. 7. In some embodiments, the counterfactual regrets for the sampled information sets in this iteration can be summed to provide aggregate value in M_R^t by information set, for example, according to Eq. (11) below and a mini-batch MCCFR algorithm as described with respect to FIG. 7. In some embodiments, duplicated records in M_S^t can be removed.

[0092] In some embodiments, if the double reservoir CFR algorithm is used, the counterfactual regret and the numerator of average strategy for the sampled information sets in this iteration (such as those saved in memories M_R^t and M_S^t in the double incremental CFR algorithm) can be saved in reservoirs M_R and M_S respectively. Reservoir sampling can be used if one or both of the reservoirs M_R and M_S is full.

[0093] As shown in lines 13-15 of the pseudocode 500, these counterfactual regret and the numerator of average strategy for the sampled information sets in this iteration can be used by a NeuralAgent algorithm as shown in FIG. 6 to optimize the two neural networks, RSN and ASN and return the parameters (e.g., θ_R^t and θ_S^t) of the RSN and ASN (e.g., $\mathcal{R}(a, I_i | \theta_R^t)$ and $S(a, I_i | \theta_S^t)$).

[0094] FIG. 6 is a pseudocode 600 of an example of an algorithm for optimizing a neural network in connection with a double neural CFR algorithm in accordance with embodiments of this specification. The example of the algorithm is referred to as NeuralAgent algorithm. The described double neural CFR algorithm can use other algorithms for optimizing the one or both neural networks used in the double neural CFR algorithm.

[0095] Define β_{epoch} as training epoch, β_{lr} as a learning rate, β_{loss} as criteria for early stopping or termination, β_{re} as the upper bound for the number of iterations from getting the minimal loss last time, θ^{t-1} as the parameter to optimize, $f(\cdot|\theta^{t-1})$ as the neural network, M as the training sample consisting information set and the corresponding target. To simplify notations, use β^* to denote the set of parameters of a neural network. For example, $\beta_{\mathcal{R}}^*$ and $\beta_{\mathcal{S}}^*$ refer to the sets of parameters in RSN and ASN, respectively. Experiments shows that the carefully designed NeuralAgent algorithm can obtain a relatively higher convergence rate of exploitability in optimizing the neural networks RSN and ASN. The pseudocode 600 shows the details of the NeuralAgent algorithm with explanatory comments.

[0096] In some embodiments, existing optimizers may not return a relatively low enough loss because of potential saddle point or local minima. To obtain a relatively higher accuracy and lower optimization loss, a scheduler is specifically designed to reduce the learning rate when the loss has stopped decrease. Specifically, the scheduler reads a metrics quantity, *e.g.* mean squared error, and if no improvement is seen for a number of epochs, the learning rate is reduced by a factor. In addition, the learning rate can be reset in both optimizer and scheduler once loss stops decrease in β_{re} epochs. Gradient clipping mechanism can be used to limit the magnitude of the parameter gradient and make optimizer behave better in the vicinity of steep cliffs. After each epoch, the best parameter will be updated. Early stopping mechanism is used once the lowest loss is less than the specified criteria β_{loss} .

[0097] In experiments, hyperparameters of the neural network can be set as follow. For example, for RSN, the neural batch size is 256 and learning rate $\beta_{lr} = 0.001$. A scheduler, who will reduce the learning rate based on the number of epochs and the convergence rate of loss, help the neural agent to obtain a high accuracy. The learning rate can be reduced by 0.5 when loss has stopped improving after 10 epochs. The lower bound on the learning rate of all parameters in this scheduler is 10^{-6} . To avoid the algorithm converging to potential local minima or saddle point, the learning rate can be reset to, for example, 0.001 and help the optimizer to learn a better performance. θ_{best}^T is the best parameters to achieve the lowest loss after T epochs. If average loss for epoch t is less than the specified criteria $\beta_{loss} = 10^{-4}$, the optimizer can have an early stop. As an example, set $\beta_{epoch} = 2000$ and update the optimizer 2000 maximum epochs.

[0098] For ASN, the loss of early stopping criteria can be set as 10^{-5} . The learning rate can be reduced by 0.7 when loss has stopped improving after 15 epochs. Other hyperparameters in ASN can be similar to those of RSN.

[0099] FIG. 7 is a pseudocode 700 of an example of a mini-batch MCCFR algorithm in accordance with embodiments of this specification. The mini-batch MCCFR algorithm (denoted as

Mini-Batch-MCCFR-NN) includes a sampling algorithm for obtaining counterfactual regret and numerator of average strategy for sampled information sets of a game. Unlike traditional outcome sampling and external sampling that only sample one block in an iteration and provide an unbiased estimator of origin CFV, the mini-batch sampling technique can randomly sample b blocks in one iteration. The example of the mini-batch MCCFR algorithm shown in the pseudocode 700 is based on the robust sampling described above. In some embodiments, the mini-batch MCCFR algorithm can be used in connection with other sampling schemes such as the depth-dependent sampling scheme. Note that the mini-batch MCCFR algorithm is an example of an algorithm for obtaining counterfactual regret and numerator of average strategy for sampled information sets of a game. The double neural CFR algorithm can use other algorithms for obtaining counterfactual regret and numerator of average strategy for sampled information sets of a game.

[0100] Let Q^j denote a block of terminals sampled according to the robust sampling scheme at j -th time, then mini-batch CFV with b mini-batches for information set I_i can be defined as:

$$\tilde{v}_i^\sigma(I_i|b) = \frac{1}{b} \sum_{j=1}^b \left(\sum_{h \in I_i, z \in Q^j, h \sqsubseteq z} \frac{\pi_{-i}^\sigma(z) \pi_i^\sigma(h, z) u_i(z)}{q(z)} \right) = \sum_{j=1}^b \frac{\tilde{v}_i^\sigma(I_i|Q^j)}{b}. \quad (11)$$

[0101] Furthermore, it can be shown that $\tilde{v}_i^\sigma(I_i|b)$ is an unbiased estimator of the counterfactual value of I_i : $E_{Q^j \sim \text{Robust Sampling}}[\tilde{v}_i^\sigma(I_i|b)] = v_i^\sigma(I_i)$.

[0102] Similarly, the cumulative mini-batch regret of action a is

$$\tilde{R}_i^T((a|I_i)|b) = \tilde{R}_i^{T-1}((a|I_i)|b) + \tilde{v}_i^{\sigma^T}((a|I_i)|b) - \tilde{v}_i^{\sigma^T}(I_i|b), \quad (12)$$

where $\tilde{R}_i^0((a|I_i)|b) = 0$. In some embodiments, the mini-batch MCCFR can sample b blocks in parallel and help MCCFR to converge faster.

[0103] Note that the mini-batch MCCFR using the regret-matching algorithm to update the cumulative mini-batch regret $\tilde{R}_i^{T,+}((a|I_i)|b)$. In some embodiments, as a variant of the Mini-Batch MCCFR, a mini-batch MCCFR+ algorithm can be used to update cumulative mini-batch regret $\tilde{R}_i^{T,+}((a|I_i)|b)$ up to iteration T by:

$$\tilde{R}_i^{T,+}((a|I_i)|b) = \begin{cases} (\tilde{v}_i^{\sigma^T}((a|I_i)|b) - \tilde{v}_i^{\sigma^T}(I_i|b))^+ & \text{if } T = 0 \\ (\tilde{R}_i^{T-1,+}((a|I_i)|b) + \tilde{v}_i^{\sigma^T}((a|I_i)|b) - \tilde{v}_i^{\sigma^T}(I_i|b))^+ & \text{if } T > 0 \end{cases} \quad (13)$$

where $(x)^+ = \max(x, 0)$. In some embodiments, it is found that mini-batch MCCFR+ converges faster than mini-batch MCCFR when specifying a suitable mini-batch size.

[0104] The function Mini-Batch-MCCFR-NN shown in the pseudocode 700 presents a mini-batch sampling method, where b blocks will be sampled in parallel. This mini-batch method can help the MCCFR to achieve a more accurate estimation of CFV. The parallel sampling makes this method efficient in practice.

[0105] As shown in lines 1-6 of the pseudocode 700, the Mini-Batch-MCCFR-NN is an iterative algorithm with an input of the total number of iterations, t . Within each iteration, an MCCFR-NN function is called for player 0 and player 1 (as shown in lines 4 and 5) and the counterfactual regret and numerator of average strategy for sampled information sets in this iteration are returned and save in memories M_R^t and M_S^t , respectively.

[0106] The function MCCFR-NN can be defined as shown in lines 8-33 of the pseudocode 700. The function MCCFR-NN traverses the game tree like tabular MCCFR, which starts from the root history $h = \emptyset$. Define I_i as the information set of h . Suppose that player i will sample k actions according to the robust sampling. Then the function can be defined as follows. (1) If the history is terminal (e.g., $h \in Z$), the function returns the weighted utility. (2) If the history is the chance player (e.g., $P(I_i) = -1$), one action $a \in A(I_i)$ can be sampled according to the strategy $\sigma_{-i}(I_i)$. Then this action will be added to the history, i.e., $h \leftarrow ha$. (3) If $P(I_i) = i$, the current strategy can be updated by the cumulative regret predicted by RSN. Then sample k actions according the specified sampling strategy profile $\sigma_i^{rs(k)}$ (e.g., robust sampling with or without depth-dependent sampling). After a recursive updating, the counterfactual value and regret of each action at I_i can be obtained. For the visited node, their counterfactual regrets and numerators of the corresponding average strategy can be stored in M_R^t and M_S^t respectively. (4) If $P(I_i)$ is the opponent, only one action will be sampled according the strategy $\sigma_{-i}(I_i)$.

[0107] FIG. 8 is a flowchart of an example of a sampling process 800 for performing MCCFR in accordance with embodiments of this specification. The sampling process 800 can be an example of the depth-dependent sampling scheme described above for performing counterfactual regret minimization (CRF) for strategy searching in strategic interaction between two or more parties. In some embodiments, strategic interaction between two or more players can be modeled by an imperfect information game (IIG) that involves two or more players. The IIG can represent one or more real-world scenarios such as resource allocation, product/service recommendation, cyber-attack prediction and/or prevention, traffic routing, fraud management, etc. that involves two or more parties, where each party may have incomplete or imperfect information about the other party's decisions. As an example, the IIG can represent a collaborative product-service recommendation service that involves at least a first player and a second player. The first player may be, for example,

an online retailer that has customer (or user) information, product and service information, purchase history of the customers, etc. The second player can be, for example, a social network platform that has social networking data of the customers, a bank or another financial institution that has financial information of the customers, a car dealership, or any other parties that may have information of the customers on the customers' preferences, needs, financial situations, locations, etc. in predicting and recommendations of products and services to the customers. The first player and the second player may each have proprietary data that does not want to share with others. The second player may only provide partial information to the first player at different times. As such, the first player may only have limited access to information of the second player. For convenience, the process 800 will be described as being performed by a data processing apparatus such as a system of one or more computers, located in one or more locations, and programmed appropriately in accordance with this specification. For example, a computer system 1000 of FIG. 10, appropriately programmed, can perform the process 800.

[0108] At 810, a data processing apparatus identifies $N1$ possible actions of a first player in a first state of the first player. In some embodiments, the IIG can be represented by a game tree (e.g., the game tree 102, 104, 202, 204 or 206). The first state of the first player can be represented by a first node of the game tree (e.g., node $h1$ of player 0 in the game tree 102) and the $N1$ possible action can be the edges or branches of the first node of the game tree (e.g., the P and B edges of the node $h1$ of player 0 in the game tree 102). In the example of the collaborative product-service recommendation service, the first state of the first player includes a history of information provided by the second player, and the $N1$ possible actions of the first player includes $N1$ possible actions in response to the history of information provided by the second player for providing product-service recommendations to customers. The first state of the first player and the possible actions can include other features in other real-world scenarios that are modeled by the IIG.

[0109] At 820, the data processing apparatus samples a possible action out of the $N1$ possible actions in the first state of the first player with a first sampling probability. In some embodiments, the data processing apparatus may sample $k1$ possible actions out of the $N1$ possible actions in the first state of the first player, wherein each of the $k1$ possible actions is sampled with the same first sampling probability.

[0110] At 830, the data processing apparatus identifies $N2$ possible actions of the first player in a second state of the first player, wherein the first state of the first player is closer to a beginning state of the IIG than the second state of the first player. In the example of the game tree 102, the second state of the first player can be, for example, the $h7$ node, which is further from the beginning state

(e.g., the h0 node) of the game tree 102 than the first state of the first player (e.g., node h1 of player 0 in the game tree 102).

[0111] At 840, the data processing apparatus samples a possible action out of the N2 possible actions in the second state of the first player with a second sampling probability, wherein the first sampling probability is less than the second sampling probability. In some embodiments, the data processing apparatus samples k2 possible actions out of the N2 possible actions in the second state of the first player, wherein each of the k2 possible actions is sampled with the same second sampling probability.

[0112] At 850, the data processing apparatus perform CRF based on the possible actions out of the N1 possible actions in the first state of the first player and the possible action out of the N2 possible actions in the second state of the first player. In some embodiments, the CRF can be performed according to the example techniques described with respect to FIG. 3 and/or FIG. 7.

[0113] In some embodiments, a strategy of the first player resulting from solving the IIG is output. The strategy can include a series of actions of the first player in the real-world scenario modeled by the IIG. For example, in the collaborative product-service recommendation scenario, the strategy of the first player resulting from solving the IIG can include, for example, a series of actions in response to the information provided by the second player, corresponding product-service recommendations to customers based on the information of the first player and the information provided by the second player. The output strategy of the first player resulting from solving the IIG can include other information in other real-world scenarios that are modeled by the IIG.

[0114] In some embodiments, performing CRF based on the possible action out of the N1 possible actions in the first state of the first player and the possible action out of the N2 possible actions in the second state of the first player includes calculating a regret value of the possible action out of the N1 possible actions in the first state of the first player (e.g., according to Eq. (1a) and/or Eq. (2)); calculating a regret value of the possible action out of the N2 possible actions in the second state of the first player (e.g., according to Eq. (1a) and/or Eq. (2)); updating a first strategy of the first player in the first state based on the regret value of the possible action out of the N1 possible actions (e.g., according to Eq. (3)); and updating a second strategy of the first player in the second state based on the regret value of the possible action out of the N2 possible actions (e.g., according to Eq. (3)).

[0115] In some embodiments, the data processing apparatus performs CRF based on the k1 possible actions out of the N1 possible actions in the first state of the first player and the k2 possible actions out of the N2 possible actions in the second state of the first player.

[0116] In some embodiments, robust sampling can be performed in connection with the depth-dependent sampling. For example, the first sampling probability is k_1/N_1 , and the second sampling probability is k_2/N_2 . As such, the possible actions are sampled according to a uniform distribution.

[0117] In some embodiments, $2 \leq k_1 \leq N_1$ and $2 \leq k_2 \leq N_2$ so that more than one possible action is visited for each state of the player.

[0118] In some embodiments, $k_1 = k_2$ so an equal number of samples are selected or visited in the first state and second state of the first player.

[0119] Similarly, the depth-dependent sampling can be performed in connection with a second player. For example, the data processing apparatus identifies M_1 possible actions of a second player in a first state of the second player. The data processing apparatus samples a possible action out of the M_1 possible actions in the first state of the second player with a third sampling probability. The data processing apparatus identifies M_2 possible actions of the second player in a second state of the second player, wherein the first state of the second player is closer to a beginning state of the IIG than the second state of the second player. The data processing apparatus samples a possible action out of the M_2 possible actions in the second state of the first player with a fourth sampling probability, wherein the third sampling probability is less than the fourth sampling probability.

[0120] In some embodiments, the depth-dependent sampling can be performed in connection with both the first player and the second player. In some embodiments, the data processing apparatus identifies M_1 possible actions of a second player in a first state of the second player, wherein the first state of the first player (e.g., state h_1 of player 0 in the game tree 102) is closer to a beginning state (e.g., h_0 state) of the IIG than the first state of the second player (e.g., state h_4 of player 1 in the game tree 102). The data processing apparatus samples a possible action out of the M_1 possible actions in the first state of the second player with a third sampling probability, wherein the third sampling probability is larger than the first sampling probability.

[0121] FIG. 9 is a flowchart of an example of double neural CFR algorithm 900 for performing MCCFR in accordance with embodiments of this specification. The sampling process 900 can be an example of the double reservoir CFR algorithm described above with respect to FIGS. 4-7 for performing counterfactual regret minimization (CRF) for strategy searching in strategic interaction between two or more players. In some embodiments, strategic interaction between two or more players can be modeled by an imperfect information game (IIG) that involves two or more players. The IIG can represent one or more real-world scenarios such as resource allocation, product/service recommendation, cyber-attack prediction and/or prevention, traffic routing, fraud management, etc. that involves two or more parties, where each party may have incomplete or imperfect information about the other party's decisions. As an example, the IIG can represent a collaborative product-

service recommendation service that involves at least a first player and a second player. The first player may be, for example, an online retailer that has customer (or user) information, product and service information, purchase history of the customers, etc. The second player can be, for example, a social network platform that has social networking data of the customers, a bank or another financial institution that has financial information of the customers, a car dealership, or any other parties that may have information of the customers on the customers' preferences, needs, financial situations, locations, etc. in predicting and recommendations of products and services to the customers. The first player and the second player may each have proprietary data that does not want to share with others. The second player may only provide partial information to the first player at different times. As such, the first player may only have limit access to information of the second player. For convenience, the process 900 will be described as being performed by a data processing apparatus such as a system of one or more computers, located in one or more locations, and programmed appropriately in accordance with this specification. For example, a computer system 1000 of FIG. 10, appropriately programmed, can perform the process 900.

[0122] At 910, a data processing apparatus initializes parameters of a first neural network and parameters of a second neural network. The first neural network (e.g., RegretSumNetwork (RSN) 420) can be used for predicting a regret value of a possible action in a state of the player. In the example of the collaborative product-service recommendation service, the state of the player includes a history of information provided by the second player, and the possible action of the player includes a possible action in response to the history of information provided by the second player for providing product-service recommendations to customers. The second neural network (e.g., AveStrategyNetwork (ASN) 430) can be used for predicting a strategy value of a possible action in a state of the player. In some embodiments, the data processing apparatus initializes parameters according to a warm start, for example, based on parameters of the first neural network and parameters of the second neural network in a previous iteration or obtained based on an existing CRF algorithm, respectively. In some embodiments, the data processing apparatus initializes parameters of the first neural network and parameters of the second neural network randomly.

[0123] At 920, the data processing apparatus stores a number of regret samples in a first data store (e.g., the reservoir M_R), wherein each of the number of regret samples includes a state of the player and a regret value of a possible action in the state of the player. In some embodiments, the regret value of a possible action in the state of the player includes a counterfactual regret value of the possible action in the state of the player calculated based on a counterfactual value of the possible action in the state of the player. For example, each regret sample can include a regret tuple $(I_i, \tilde{r}_i^{\sigma_i}(a|I_i))$. In some embodiments, the number of regret samples are obtained in two or more

iterations of a counterfactual regret minimization (CRF) algorithm in strategy searching in strategic interaction between the player and at least another player. In some embodiments, the CRF algorithm includes a robust sampling CRF algorithm.

[0124] In some embodiments, in each of the two or more iterations of the CRF algorithm in solving the IIG, the data processing apparatus samples a possible action out of a number of possible actions in a second state of a player according to a sampling scheme; calculates a counterfactual value of the possible action in the second state of the player (e.g., according to Eq. (1)); calculates a regret value of the possible action in the second state of the player based on the counterfactual value of the possible action in the second state of the player (e.g., according to Eq. (1a) and/or Eq. (2)); calculate an updated strategy of the possible action in the second state of the player based the regret value of the possible action in the second state of the player according to a regret matching algorithm (e.g., according to Eq. (3)); and calculates a strategy value of the possible action in the second state of the player based on the updated strategy of the possible action in the second state of the player (e.g., according to Eq. (4) and/or Eq. (5)).

[0125] In some embodiments, the data processing apparatus can obtain a new regret sample (e.g., by performing another iteration of the MCCFR). The data processing apparatus can store the new regret sample into the first data store according to a reservoir sampling algorithm. For example, storing the new regret sample into the first data store according to a reservoir sampling algorithm includes: determining whether the first data store is full; and in response to determining that the first data store is full, replacing one of the number of regret samples in the first data store with the new regret sample.

[0126] At 930, the data processing apparatus stores a number of strategy samples in a second data store (e.g., the reservoir M_S), wherein each of the number of strategy samples includes a state of the player and a strategy value of a possible action in the state of the player. In some embodiments, the strategy value of a possible action in the state of the player includes a numerator of an average strategy. For example, each of the number of strategy samples can include a strategy tuple $(I_i, s^t(a|I_i))$.

[0127] At 940, the data processing apparatus updates parameters of a first neural network for predicting a regret value of a possible action in a state of the player based on the number of regret samples in the first data store, for example, according to Eq. (7). In some embodiments, the parameters of the first neural network can be updated according to the NeuralAgent algorithm shown in FIG. 6, or any other algorithms for optimizing a neural network.

[0128] At 950, the data processing apparatus updates parameters of a second neural network for predicting a strategy value of a possible action in a state of the player based on the number of strategy samples in the second data store, for example, according to Eq. (8). In some embodiments, the parameters of the second neural network can be updated according to the NeuralAgent algorithm shown in FIG. 6, or any other algorithms for optimizing a neural network.

[0129] At 960, the data processing apparatus identifies a first state of the player and a first possible action in the first state of the player.

[0130] At 970, the data processing apparatus predicts a first regret value of the first possible action in the first state of the player using the parameters of the first neural network. In some embodiments, the predicted first regret value of the first possible action in the first state of the player can be used in a next iteration of the CFR algorithm.

[0131] At 980, the data processing apparatus predicts a first strategy value of the first possible action in the first state of the player using the parameters of the second neural network. In some embodiments, the predicted first strategy value of the first possible action in the first state of the player can be used in a next iteration of the CFR algorithm. In some embodiments, the predicted first strategy value of the first possible action in the first state of the player can be used to calculate an approximate Nash equilibrium and serve as an output of the CFR algorithm. In some embodiments, the predicted first strategy value of the first possible action in the first state of the player can include a series of actions of the first player in the real-world scenario modeled by the IIG. For example, in the collaborative product-service recommendation scenario, the predicted first strategy value of the first possible action in the first state of the player can include, for example, a series of actions in response to the information provided by the second player, corresponding product-service recommendations to customers based on the information of the first player and the information provided by the second player. The predicted first strategy value of the first possible action in the first state of the player can include other information in other real-world scenarios that are modeled by the IIG.

[0132] FIG. 10 depicts a block diagram illustrating an example of a computer-implemented system used to provide computational functionalities associated with described algorithms, methods, functions, processes, flows, and procedures in accordance with embodiments of this specification. FIG. 10 is a block diagram illustrating an example of a computer-implemented System 1000 used to provide computational functionalities associated with described algorithms, methods, functions, processes, flows, and procedures, according to an embodiment of the present disclosure. In the illustrated embodiment, System 1000 includes a Computer 1002 and a Network 1030.

[0133] The illustrated Computer 1002 is intended to encompass any computing device such as a server, desktop computer, laptop/notebook computer, wireless data port, smart phone, personal data assistant (PDA), tablet computer, one or more processors within these devices, another computing device, or a combination of computing devices, including physical or virtual instances of the computing device, or a combination of physical or virtual instances of the computing device. Additionally, the Computer 1002 can include an input device, such as a keypad, keyboard, touch screen, another input device, or a combination of input devices that can accept user information, and an output device that conveys information associated with the operation of the Computer 1002, including digital data, visual, audio, another type of information, or a combination of types of information, on a graphical-type user interface (UI) (or GUI) or other UI.

[0134] The Computer 1002 can serve in a role in a distributed computing system as a client, network component, a server, a database or another persistency, another role, or a combination of roles for performing the subject matter described in the present disclosure. The illustrated Computer 1002 is communicably coupled with a Network 1030. In some embodiments, one or more components of the Computer 1002 can be configured to operate within an environment, including cloud-computing-based, local, global, another environment, or a combination of environments.

[0135] At a high level, the Computer 1002 is an electronic computing device operable to receive, transmit, process, store, or manage data and information associated with the described subject matter. According to some embodiments, the Computer 1002 can also include or be communicably coupled with a server, including an application server, e-mail server, web server, caching server, streaming data server, another server, or a combination of servers.

[0136] The Computer 1002 can receive requests over Network 1030 (for example, from a client software application executing on another Computer 1002) and respond to the received requests by processing the received requests using a software application or a combination of software applications. In addition, requests can also be sent to the Computer 1002 from internal users (for example, from a command console or by another internal access method), external or third-parties, or other entities, individuals, systems, or computers.

[0137] Each of the components of the Computer 1002 can communicate using a System Bus 1003. In some embodiments, any or all of the components of the Computer 1002, including hardware, software, or a combination of hardware and software, can interface over the System Bus 1003 using an application programming interface (API) 1012, a Service Layer 1013, or a combination of the API 1012 and Service Layer 1013. The API 1012 can include specifications for routines, data structures, and object classes. The API 1012 can be either computer-language independent or dependent and refer to a complete interface, a single function, or even a set of APIs.

The Service Layer 1013 provides software services to the Computer 1002 or other components (whether illustrated or not) that are communicably coupled to the Computer 1002. The functionality of the Computer 1002 can be accessible for all service consumers using the Service Layer 1013. Software services, such as those provided by the Service Layer 1013, provide reusable, defined functionalities through a defined interface. For example, the interface can be software written in JAVA, C++, another computing language, or a combination of computing languages providing data in extensible markup language (XML) format, another format, or a combination of formats. While illustrated as an integrated component of the Computer 1002, alternative embodiments can illustrate the API 1012 or the Service Layer 1013 as stand-alone components in relation to other components of the Computer 1002 or other components (whether illustrated or not) that are communicably coupled to the Computer 1002. Moreover, any or all parts of the API 1012 or the Service Layer 1013 can be implemented as a child or a sub-module of another software module, enterprise application, or hardware module without departing from the scope of the present disclosure.

[0138] The Computer 1002 includes an Interface 1004. Although illustrated as a single Interface 1004, two or more Interfaces 1004 can be used according to particular needs, desires, or particular embodiments of the Computer 1002. The Interface 1004 is used by the Computer 1002 for communicating with another computing system (whether illustrated or not) that is communicatively linked to the Network 1030 in a distributed environment. Generally, the Interface 1004 is operable to communicate with the Network 1030 and includes logic encoded in software, hardware, or a combination of software and hardware. More specifically, the Interface 1004 can include software supporting one or more communication protocols associated with communications such that the Network 1030 or hardware of Interface 1004 is operable to communicate physical signals within and outside of the illustrated Computer 1002.

[0139] The Computer 1002 includes a Processor 1005. Although illustrated as a single Processor 1005, two or more Processors 1005 can be used according to particular needs, desires, or particular embodiments of the Computer 1002. Generally, the Processor 1005 executes instructions and manipulates data to perform the operations of the Computer 1002 and any algorithms, methods, functions, processes, flows, and procedures as described in the present disclosure.

[0140] The Computer 1002 also includes a Database 1006 that can hold data for the Computer 1002, another component communicatively linked to the Network 1030 (whether illustrated or not), or a combination of the Computer 1002 and another component. For example, Database 1006 can be an in-memory, conventional, or another type of database storing data consistent with the present disclosure. In some embodiments, Database 1006 can be a combination of two or more different database types (for example, a hybrid in-memory and conventional database) according to particular

needs, desires, or particular embodiments of the Computer 1002 and the described functionality. Although illustrated as a single Database 1006, two or more databases of similar or differing types can be used according to particular needs, desires, or particular embodiments of the Computer 1002 and the described functionality. While Database 1006 is illustrated as an integral component of the Computer 1002, in alternative embodiments, Database 1006 can be external to the Computer 1002. As an example, Database 1006 can include the above-described reservoir M_R 1016 that store regret samples 1026 and reservoir M_S 1018 that store strategy samples 1028.

[0141] The Computer 1002 also includes a Memory 1007 that can hold data for the Computer 1002, another component or components communicatively linked to the Network 1030 (whether illustrated or not), or a combination of the Computer 1002 and another component. Memory 1007 can store any data consistent with the present disclosure. In some embodiments, Memory 1007 can be a combination of two or more different types of memory (for example, a combination of semiconductor and magnetic storage) according to particular needs, desires, or particular embodiments of the Computer 1002 and the described functionality. Although illustrated as a single Memory 1007, two or more Memories 1007 or similar or differing types can be used according to particular needs, desires, or particular embodiments of the Computer 1002 and the described functionality. While Memory 1007 is illustrated as an integral component of the Computer 1002, in alternative embodiments, Memory 1007 can be external to the Computer 1002.

[0142] The Application 1008 is an algorithmic software engine providing functionality according to particular needs, desires, or particular embodiments of the Computer 1002, particularly with respect to functionality described in the present disclosure. For example, Application 1008 can serve as one or more components, modules, or applications. Further, although illustrated as a single Application 1008, the Application 1008 can be implemented as multiple Applications 1008 on the Computer 1002. In addition, although illustrated as integral to the Computer 1002, in alternative embodiments, the Application 1008 can be external to the Computer 1002.

[0143] The Computer 1002 can also include a Power Supply 1014. The Power Supply 1014 can include a rechargeable or non-rechargeable battery that can be configured to be either user- or non-user-replaceable. In some embodiments, the Power Supply 1014 can include power-conversion or management circuits (including recharging, standby, or another power management functionality). In some embodiments, the Power Supply 1014 can include a power plug to allow the Computer 1002 to be plugged into a wall socket or another power source to, for example, power the Computer 1002 or recharge a rechargeable battery.

[0144] There can be any number of Computers 1002 associated with, or external to, a computer system containing Computer 1002, each Computer 1002 communicating over Network 1030.

Further, the term “client,” “user,” or other appropriate terminology can be used interchangeably, as appropriate, without departing from the scope of the present disclosure. Moreover, the present disclosure contemplates that many users can use one Computer 1002, or that one user can use multiple computers 1002.

[0145] FIG. 11 is a diagram of an example of modules of an apparatus 1100 in accordance with embodiments of this specification. The apparatus 1100 can be an example embodiment of a data processing apparatus for performing counterfactual regret minimization (CRF) for strategy searching in strategic interaction between two or more players. In some embodiments, strategic interaction between two or more players can be modeled by an imperfect information game (IIG) that involves two or more players. As an example, the IIG represents a collaborative product-service recommendation service that involves at least a first player and a second player, the first player having limited access to information of the second player. The apparatus 1100 can correspond to the embodiments described above, and the apparatus 1100 includes the following: a first identifying module 1101 for identifying N1 possible actions of a first player in a first state of the first player; a first sampling module 1102 for sampling a possible action out of the N1 possible actions in the first state of the first player with a first sampling probability; a second identifying module 1103 for identifying N2 possible actions of the first player in a second state of the first player, wherein the first state of the first player is closer to a beginning state of the IIG than the second state of the first player; a second sampling module 1104 for sampling a possible action out of the N2 possible actions in the second state of the first player with a second sampling probability, wherein the first sampling probability is less than the second sampling probability; and a processing module 1105 for performing the CRF based on the possible actions out of the N1 possible actions in the first state of the first player and the possible action out of the N2 possible actions in the second state of the first player. In some embodiments, the first state of the first player includes a history of information provided by the second player, and the N1 possible actions of the first player includes N1 possible actions in response to the history of information provided by the second player for providing product-service recommendations to customers.

[0146] In an optional embodiment, the processing module includes: a first calculating module for calculating a regret value of the possible action out of the N1 possible actions in the first state of the first player; a second calculating module for calculating a regret value of the possible action out of the N2 possible actions in the second state of the first player; a first updating module for updating a first strategy of the first player in the first state based on the regret value of the possible action out of the N1 possible actions; and a second updating module for updating a second strategy of the first

player in the second state based on the regret value of the possible action out of the N2 possible actions.

[0147] In an optional embodiment, the first sampling module samples k1 possible actions out of the N1 possible actions in the first state of the first player, wherein each of the k1 possible actions is sampled with the same first sampling probability; and the second sampling module samples k2 possible actions out of the N2 possible actions in the second state of the first player, wherein each of the k2 possible actions is sampled with the same second sampling probability.

[0148] In an optional embodiment, the processing module performs the CRF based on the k1 possible actions out of the N1 possible actions in the first state of the first player and the k2 possible actions out of the N2 possible actions in the second state of the first player.

[0149] In an optional embodiment, the first sampling probability is $k1/N1$, and the second sampling probability is $k2/N2$.

[0150] In an optional embodiment, $2 \leq k1 \leq N1$ and $2 \leq k2 \leq N2$.

[0151] In an optional embodiment, $k1 = k2$.

[0152] In an optional embodiment, the apparatus 1100 further includes the following: a third identifying module for identifying M1 possible actions of a second player in a first state of the second player; a third sampling module for sampling a possible action out of the M1 possible actions in the first state of the second player with a third sampling probability; a fourth identifying module for identifying M2 possible actions of the second player in a second state of the second player, wherein the first state of the second player is closer to a beginning state of the IIG than the second state of the second player; and a fourth sampling module for sampling a possible action out of the M2 possible actions in the second state of the first player with a fourth sampling probability, wherein the third sampling probability is less than the fourth sampling probability.

[0153] In an optional embodiment, the apparatus 1100 further includes the following: a fifth identifying module for identifying M1 possible actions of a second player in a first state of the second player, wherein the first state of the first player is closer to a beginning state of the IIG than the first state of the second player; and a fifth sampling module for sampling a possible action out of the M1 possible actions in the first state of the second player with a third sampling probability, wherein the third sampling probability is larger than the first sampling probability.

[0154] The system, apparatus, module, or unit illustrated in the previous embodiments can be implemented by using a computer chip or an entity, or can be implemented by using a product having a certain function. A typical embodiment device is a computer, and the computer can be a personal computer, a laptop computer, a cellular phone, a camera phone, a smartphone, a personal

digital assistant, a media player, a navigation device, an email receiving and sending device, a game console, a tablet computer, a wearable device, or any combination of these devices.

[0155] For an embodiment process of functions and roles of each module in the apparatus, references can be made to an embodiment process of corresponding steps in the previous method. Details are omitted here for simplicity.

[0156] Because an apparatus embodiment basically corresponds to a method embodiment, for related parts, references can be made to related descriptions in the method embodiment. The previously described apparatus embodiment is merely an example. The modules described as separate parts may or may not be physically separate, and parts displayed as modules may or may not be physical modules, may be located in one position, or may be distributed on a number of network modules. Some or all of the modules can be selected based on actual demands to achieve the objectives of the solutions of the specification. A person of ordinary skill in the art can understand and implement the embodiments of the present application without creative efforts.

[0157] Referring again to FIG. 11, it can be interpreted as illustrating an internal functional module and a structure of a data processing apparatus for performing counterfactual regret minimization (CRF) for strategy searching in strategic interaction between two or more players. In some embodiments, strategic interaction between two or more players can be modeled by an imperfect information game (IIG) that involves two or more players. An execution body in essence can be an electronic device, and the electronic device includes the following: one or more processors; and a memory configured to store an executable instruction of the one or more processors.

[0158] The one or more processors are configured to identify N1 possible actions of a first player in a first state of the first player; sample a possible action out of the N1 possible actions in the first state of the first player with a first sampling probability; identify N2 possible actions of the first player in a second state of the first player, wherein the first state of the first player is closer to a beginning state of the IIG than the second state of the first player; sample a possible action out of the N2 possible actions in the second state of the first player with a second sampling probability, wherein the first sampling probability is less than the second sampling probability; and perform the CRF based on the possible actions out of the N1 possible actions in the first state of the first player and the possible action out of the N2 possible actions in the second state of the first player.

[0159] Optionally, the one or more processors are configured to calculate a regret value of the possible action out of the N1 possible actions in the first state of the first player; calculate a regret value of the possible action out of the N2 possible actions in the second state of the first player; update a first strategy of the first player in the first state based on the regret value of the possible

action out of the N_1 possible actions; and update a second strategy of the first player in the second state based on the regret value of the possible action out of the N_2 possible actions.

[0160] Optionally, the one or more processors are configured to sample k_1 possible actions out of the N_1 possible actions in the first state of the first player, wherein each of the k_1 possible actions is sampled with the same first sampling probability; and sample k_2 possible actions out of the N_2 possible actions in the second state of the first player, wherein each of the k_2 possible actions is sampled with the same second sampling probability.

[0161] Optionally, the one or more processors are configured to perform CRF based on the k_1 possible actions out of the N_1 possible actions in the first state of the first player and the k_2 possible actions out of the N_2 possible actions in the second state of the first player.

[0162] Optionally, the first sampling probability is k_1/N_1 , and the second sampling probability is k_2/N_2 .

[0163] Optionally, $2 \leq k_1 \leq N_1$ and $2 \leq k_2 \leq N_2$.

[0164] Optionally, $k_1 = k_2$.

[0165] Optionally, the one or more processors are configured to identify M_1 possible actions of a second player in a first state of the second player; sample a possible action out of the M_1 possible actions in the first state of the second player with a third sampling probability; identify M_2 possible actions of the second player in a second state of the second player, wherein the first state of the second player is closer to a beginning state of the IIG than the second state of the second player; and sample a possible action out of the M_2 possible actions in the second state of the first player with a fourth sampling probability, wherein the third sampling probability is less than the fourth sampling probability.

[0166] Optionally, the one or more processors are configured to: identify M_1 possible actions of a second player in a first state of the second player, wherein the first state of the first player is closer to a beginning state of the IIG than the first state of the second player; and sample a possible action out of the M_1 possible actions in the first state of the second player with a third sampling probability, wherein the third sampling probability is larger than the first sampling probability.

[0167] FIG. 12 is a diagram of an example of modules of another apparatus 1200 in accordance with embodiments of this specification. The apparatus 1200 can be an example embodiment of a data processing apparatus for performing counterfactual regret minimization (CRF) for strategy searching in strategic interaction between two or more players. The apparatus 1200 can correspond to the embodiments described above, and the apparatus 1200 includes the following: a first storing module 1201 for storing a number of regret samples in a first data store, wherein each of the number of regret samples includes a state of a player and a regret value of a possible action in the state of the player,

wherein the number of regret samples are obtained in two or more iterations of a counterfactual regret minimization (CRF) algorithm in strategy searching in strategic interaction between the player and at least another player; a second storing module 1202 for storing a number of strategy samples in a second data store, wherein each of the number of strategy samples includes a state of the player and a strategy value of a possible action in the state of the player; a first updating module 1203 for updating parameters of a first neural network for predicting a regret value of a possible action in a state of the player based on the number of regret samples in the first data store; and a second updating module 1204 for updating parameters of a second neural network for predicting a strategy value of a possible action in a state of the player based on the number of strategy samples in the second data store. In some embodiments, strategic interaction between two or more players can be modeled by an imperfect information game (IIG) that involves two or more players. As an example, the IIG represents a collaborative product-service recommendation service that involves at least the player and a second player, the player having limited access to information of the second player, wherein the state of the player includes a history of information provided by the second player, and wherein the possible action of the player includes a possible action in response to the history of information provided by the second player for providing product-service recommendations to customers.

[0168] The foregoing and other described embodiments can each, optionally, include one or more of the following features:

[0169] In an optional embodiment, the apparatus 1200 further includes the following: an identifying module for identifying a first state of the player and a first possible action in the first state of the player; a first predicting module for predicting a first regret value of the first possible action in the first state of the player using the parameters of the first neural network; and a second predicting module for predicting a first strategy value of the first possible action in the first state of the player using the parameters of the second neural network.

[0170] In an optional embodiment, wherein the first storing module is capable of obtaining a new regret sample; and storing the new regret sample into the first data store according to a reservoir sampling algorithm.

[0171] In an optional embodiment, wherein storing the new regret sample into the first data store according to a reservoir sampling algorithm includes: determining whether the first data store is full; and in response to determining that the first data store is full, replacing one of the number of regret samples in the first data store with the new regret sample.

[0172] In an optional embodiment, wherein the CRF algorithm includes a robust sampling CRF algorithm.

[0173] In an optional embodiment, wherein the strategy value of a possible action in the state of the player includes a numerator of an average strategy.

[0174] In an optional embodiment, wherein the regret value of a possible action in the state of the player includes a counterfactual regret value of the possible action in the state of the player calculated based on a counterfactual value of the possible action in the state of the player.

[0175] In an optional embodiment, the apparatus 1200 further includes the following: further including: in each of the two or more iterations of a counterfactual regret minimization (CRF) algorithm in strategy searching in strategic interaction between the player and the at least another player, a sampling module for sampling a possible action out of a number of possible actions in a second state of a player according to a sampling scheme; a first calculating module for calculating a counterfactual value of the possible action in the second state of the player; a second calculating module for calculating a regret value of the possible action in the second state of the player based on the counterfactual value of the possible action in the second state of the player; a third calculating module for calculating an updated strategy of the possible action in the second state of the player based the regret value of the possible action in the second state of the player according to a regret matching algorithm; and a fourth calculating module for calculating a strategy value of the possible action in the second state of the player based on the updated strategy of the possible action in the second state of the player.

[0176] In an optional embodiment, the apparatus 1200 further includes the following: further including: a first initializing module for initializing the parameters of the first neural network based on parameters of the first neural network in a previous iteration; and a second initializing module for initializing the parameters of the second neural network based on parameters of the second neural network in a previous iteration.

[0177] The system, apparatus, module, or module illustrated in the previous embodiments can be implemented by using a computer chip or an entity, or can be implemented by using a product having a certain function. A typical embodiment device is a computer, and the computer can be a personal computer, a laptop computer, a cellular phone, a camera phone, a smartphone, a personal digital assistant, a media player, a navigation device, an email receiving and sending device, a game console, a tablet computer, a wearable device, or any combination of these devices.

[0178] For an embodiment process of functions and roles of each module in the apparatus, references can be made to an embodiment process of corresponding steps in the previous method. Details are omitted here for simplicity.

[0179] Because an apparatus embodiment basically corresponds to a method embodiment, for related parts, references can be made to related descriptions in the method embodiment. The

previously described apparatus embodiment is merely an example. The modules described as separate parts may or may not be physically separate, and parts displayed as modules may or may not be physical modules, may be located in one position, or may be distributed on a number of network modules. Some or all of the modules can be selected based on actual demands to achieve the objectives of the solutions of the specification. A person of ordinary skill in the art can understand and implement the embodiments of the present application without creative efforts.

[0180] Referring again to FIG. 12, it can be interpreted as illustrating an internal functional module and a structure of a data processing apparatus for performing counterfactual regret minimization (CRF) for strategy searching in strategic interaction between two or more players. An execution body in essence can be an electronic device, and the electronic device includes the following: one or more processors; and a memory configured to store an executable instruction of the one or more processors.

[0181] The one or more processors are configured to store a number of regret samples in a first data store, wherein each of the number of regret samples includes a state of a player and a regret value of a possible action in the state of the player, wherein the number of regret samples are obtained in two or more iterations of a counterfactual regret minimization (CRF) algorithm in strategy searching in strategic interaction between the player and at least another player; store a number of strategy samples in a second data store, wherein each of the number of strategy samples includes a state of the player and a strategy value of a possible action in the state of the player; update parameters of a first neural network for predicting a regret value of a possible action in a state of the player based on the number of regret samples in the first data store; and update parameters of a second neural network for predicting a strategy value of a possible action in a state of the player based on the number of strategy samples in the second data store. In some embodiments, strategic interaction between two or more players can be modeled by an imperfect information game (IIG) that involves two or more players. As an example, the IIG represents a collaborative product-service recommendation service that involves at least the player and a second player, the player having limited access to information of the second player, wherein the state of the player includes a history of information provided by the second player, and wherein the possible action of the player includes a possible action in response to the history of information provided by the second player for providing product-service recommendations to customers.

[0182] Optionally, the one or more processors are configured to: identify a first state of the player and a first possible action in the first state of the player; predict a first regret value of the first possible action in the first state of the player using the parameters of the first neural network; and

predict a first strategy value of the first possible action in the first state of the player using the parameters of the second neural network.

[0183] Optionally, the one or more processors are configured to: obtain a new regret sample; and store the new regret sample into the first data store according to a reservoir sampling algorithm.

[0184] Optionally, the one or more processors are configured to: determine whether the first data store is full; and in response to determining that the first data store is full, replace one of the number of regret samples in the first data store with the new regret sample.

[0185] Optionally, the CRF algorithm includes a robust sampling CRF algorithm.

[0186] Optionally, the strategy value of a possible action in the state of the player includes a numerator of an average strategy.

[0187] Optionally, the regret value of a possible action in the state of the player includes a counterfactual regret value of the possible action in the state of the player calculated based on a counterfactual value of the possible action in the state of the player.

[0188] Optionally, the one or more processors are configured to: in each of the two or more iterations of a counterfactual regret minimization (CRF) algorithm in strategy searching in strategic interaction between the player and the at least another player, sample a possible action out of a number of possible actions in a second state of a player according to a sampling scheme; calculate a counterfactual value of the possible action in the second state of the player; calculate a regret value of the possible action in the second state of the player based on the counterfactual value of the possible action in the second state of the player; calculate an updated strategy of the possible action in the second state of the player based the regret value of the possible action in the second state of the player according to a regret matching algorithm; and calculate a strategy value of the possible action in the second state of the player based on the updated strategy of the possible action in the second state of the player.

[0189] Optionally, the one or more processors are configured to initialize the parameters of the first neural network based on parameters of the first neural network in a previous iteration; and initialize the parameters of the second neural network based on parameters of the second neural network in a previous iteration.

[0190] Described embodiments of the subject matter can include one or more features, alone or in combination. For example, in a first embodiment, a computer-implemented method for performing counterfactual regret minimization (CRF) for strategy searching in strategic interaction between two or more players. The method includes: storing a number of regret samples in a first data store, wherein each of the number of regret samples includes a state of a player and a regret value of a possible action in the state of the player, wherein the number of regret samples are obtained in two

or more iterations of a counterfactual regret minimization (CRF) algorithm in strategy searching in strategic interaction between the player and at least another player; storing a number of strategy samples in a second data store, wherein each of the number of strategy samples includes a state of the player and a strategy value of a possible action in the state of the player; updating parameters of a first neural network for predicting a regret value of a possible action in a state of the player based on the number of regret samples in the first data store; and updating parameters of a second neural network for predicting a strategy value of a possible action in a state of the player based on the number of strategy samples in the second data store.

[0191] The foregoing and other described embodiments can each, optionally, include one or more of the following features:

[0192] A first feature, combinable with any of the following features, further including: identifying a first state of the player and a first possible action in the first state of the player; predicting a first regret value of the first possible action in the first state of the player using the parameters of the first neural network; and predicting a first strategy value of the first possible action in the first state of the player using the parameters of the second neural network.

[0193] A second feature, combinable with any of the following features, further including: obtaining a new regret sample; and storing the new regret sample into the first data store according to a reservoir sampling algorithm.

[0194] A third feature, combinable with any of the following features, wherein storing the new regret sample into the first data store according to a reservoir sampling algorithm includes: determining whether the first data store is full; and in response to determining that the first data store is full, replacing one of the number of regret samples in the first data store with the new regret sample.

[0195] A fourth feature, combinable with any of the following features, wherein the CRF algorithm includes a robust sampling CRF algorithm.

[0196] A fifth feature, combinable with any of the following features, wherein the strategy value of a possible action in the state of the player includes a numerator of an average strategy.

[0197] A sixth feature, combinable with any of the following features, wherein the regret value of a possible action in the state of the player includes a counterfactual regret value of the possible action in the state of the player calculated based on a counterfactual value of the possible action in the state of the player.

[0198] A seventh feature, combinable with any of the following features, further including: in each of the two or more iterations of a counterfactual regret minimization (CRF) algorithm in strategy searching in strategic interaction between the player and the at least another player, sampling

a possible action out of a number of possible actions in a second state of a player according to a sampling scheme; calculating a counterfactual value of the possible action in the second state of the player; calculating a regret value of the possible action in the second state of the player based on the counterfactual value of the possible action in the second state of the player; calculating an updated strategy of the possible action in the second state of the player based the regret value of the possible action in the second state of the player according to a regret matching algorithm; and calculating a strategy value of the possible action in the second state of the player based on the updated strategy of the possible action in the second state of the player.

[0199] An eighth feature, combinable with any of the following features, further including: initializing the parameters of the first neural network based on parameters of the first neural network in a previous iteration; and initializing the parameters of the second neural network based on parameters of the second neural network in a previous iteration.

[0200] In a second embodiment, a system, including: one or more processors; and one or more computer-readable memories coupled to the one or more processors and having instructions stored thereon which are executable by the one or more processors to perform the method of any of the first embodiment and its optional combination of the one or more of features described above.

[0201] In a third embodiment, an apparatus for performing counterfactual regret minimization (CRF) for strategy searching in strategic interaction involves two or more players, including: a first storing module for storing a number of regret samples in a first data store, wherein each of the number of regret samples includes a state of a player and a regret value of a possible action in the state of the player, wherein the number of regret samples are obtained in two or more iterations of a counterfactual regret minimization (CRF) algorithm in strategy searching in strategic interaction between the player and at least another player; a second storing module for storing a number of strategy samples in a second data store, wherein each of the number of strategy samples includes a state of the player and a strategy value of a possible action in the state of the player; a first updating module for updating parameters of a first neural network for predicting a regret value of a possible action in a state of the player based on the number of regret samples in the first data store; and a second updating module for updating parameters of a second neural network for predicting a strategy value of a possible action in a state of the player based on the number of strategy samples in the second data store. In some embodiments, strategic interaction between two or more players can be modeled by an imperfect information game (IIG) that involves two or more players. As an example, the IIG represents a collaborative product-service recommendation service that involves at least the player and a second player, the player having limited access to information of the second player, wherein the state of the player includes a history of information provided by the second player, and

wherein the possible action of the player includes a possible action in response to the history of information provided by the second player for providing product-service recommendations to customers.

[0202] The foregoing and other described embodiments can each, optionally, include one or more of the following features:

[0203] A first feature, combinable with any of the following features, further including: an identifying module for identifying a first state of the player and a first possible action in the first state of the player; a first predicting module for predicting a first regret value of the first possible action in the first state of the player using the parameters of the first neural network; and a second predicting module for predicting a first strategy value of the first possible action in the first state of the player using the parameters of the second neural network.

[0204] A second feature, combinable with any of the following features, wherein the first storing module is capable of obtaining a new regret sample; and storing the new regret sample into the first data store according to a reservoir sampling algorithm.

[0205] A third feature, combinable with any of the following features, wherein storing the new regret sample into the first data store according to a reservoir sampling algorithm includes: determining whether the first data store is full; and in response to determining that the first data store is full, replacing one of the number of regret samples in the first data store with the new regret sample.

[0206] A fourth feature, combinable with any of the following features, wherein the CRF algorithm includes a robust sampling CRF algorithm.

[0207] A fifth feature, combinable with any of the following features, wherein the strategy value of a possible action in the state of the player includes a numerator of an average strategy.

[0208] A sixth feature, combinable with any of the following features, wherein the regret value of a possible action in the state of the player includes a counterfactual regret value of the possible action in the state of the player calculated based on a counterfactual value of the possible action in the state of the player.

[0209] A seventh feature, combinable with any of the following features, further including: in each of the two or more iterations of a counterfactual regret minimization (CRF) algorithm in strategy searching in strategic interaction between the player and the at least another player, a sampling module for sampling a possible action out of a number of possible actions in a second state of a player according to a sampling scheme; a first calculating module for calculating a counterfactual value of the possible action in the second state of the player; a second calculating module for calculating a regret value of the possible action in the second state of the player based on

the counterfactual value of the possible action in the second state of the player; a third calculating module for calculating an updated strategy of the possible action in the second state of the player based the regret value of the possible action in the second state of the player according to a regret matching algorithm; and a fourth calculating module for calculating a strategy value of the possible action in the second state of the player based on the updated strategy of the possible action in the second state of the player.

[0210] An eighth feature, combinable with any of the following features, further including: a first initializing module for initializing the parameters of the first neural network based on parameters of the first neural network in a previous iteration; and a second initializing module for initializing the parameters of the second neural network based on parameters of the second neural network in a previous iteration.

[0211] Embodiments of the subject matter and the actions and operations described in this specification can be implemented in digital electronic circuitry, in tangibly-embodied computer software or firmware, in computer hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Embodiments of the subject matter described in this specification can be implemented as one or more computer programs, e.g., one or more modules of computer program instructions, encoded on a computer program carrier, for execution by, or to control the operation of, data processing apparatus. For example, a computer program carrier can include one or more computer-readable storage media that have instructions encoded or stored thereon. The carrier may be a tangible non-transitory computer-readable medium, such as a magnetic, magneto optical, or optical disk, a solid state drive, a random access memory (RAM), a read-only memory (ROM), or other types of media. Alternatively, or in addition, the carrier may be an artificially generated propagated signal, e.g., a machine-generated electrical, optical, or electromagnetic signal that is generated to encode information for transmission to suitable receiver apparatus for execution by a data processing apparatus. The computer storage medium can be or be part of a machine-readable storage device, a machine-readable storage substrate, a random or serial access memory device, or a combination of one or more of them. A computer storage medium is not a propagated signal.

[0212] A computer program, which may also be referred to or described as a program, software, a software application, an app, a module, a software module, an engine, a script, or code, can be written in any form of programming language, including compiled or interpreted languages, or declarative or procedural languages; and it can be deployed in any form, including as a stand-alone program or as a module, component, engine, subroutine, or other unit suitable for executing in a

computing environment, which environment may include one or more computers interconnected by a data communication network in one or more locations.

[0213] A computer program may, but need not, correspond to a file in a file system. A computer program can be stored in a portion of a file that holds other programs or data, e.g., one or more scripts stored in a markup language document, in a single file dedicated to the program in question, or in multiple coordinated files, e.g., files that store one or more modules, sub programs, or portions of code.

[0214] Processors for execution of a computer program include, by way of example, both general- and special-purpose microprocessors, and any one or more processors of any kind of digital computer. Generally, a processor will receive the instructions of the computer program for execution as well as data from a non-transitory computer-readable medium coupled to the processor.

[0215] The term “data processing apparatus” encompasses all kinds of apparatuses, devices, and machines for processing data, including by way of example a programmable processor, a computer, or multiple processors or computers. Data processing apparatus can include special-purpose logic circuitry, e.g., an FPGA (field programmable gate array), an ASIC (application specific integrated circuit), or a GPU (graphics processing unit). The apparatus can also include, in addition to hardware, code that creates an execution environment for computer programs, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them.

[0216] The processes and logic flows described in this specification can be performed by one or more computers or processors executing one or more computer programs to perform operations by operating on input data and generating output. The processes and logic flows can also be performed by special-purpose logic circuitry, e.g., an FPGA, an ASIC, or a GPU, or by a combination of special-purpose logic circuitry and one or more programmed computers.

[0217] Computers suitable for the execution of a computer program can be based on general or special-purpose microprocessors or both, or any other kind of central processing unit. Generally, a central processing unit will receive instructions and data from a read only memory or a random access memory or both. Elements of a computer can include a central processing unit for executing instructions and one or more memory devices for storing instructions and data. The central processing unit and the memory can be supplemented by, or incorporated in, special-purpose logic circuitry.

[0218] Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to one or more storage devices. The storage devices can be, for example, magnetic, magneto optical, or optical disks, solid state drives, or any other type of non-transitory, computer-

readable media. However, a computer need not have such devices. Thus, a computer may be coupled to one or more storage devices, such as, one or more memories, that are local and/or remote. For example, a computer can include one or more local memories that are integral components of the computer, or the computer can be coupled to one or more remote memories that are in a cloud network. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver, or a portable storage device, e.g., a universal serial bus (USB) flash drive, to name just a few.

[0219] Components can be “coupled to” each other by being commutatively such as electrically or optically connected to one another, either directly or via one or more intermediate components. Components can also be “coupled to” each other if one of the components is integrated into the other. For example, a storage component that is integrated into a processor (e.g., an L2 cache component) is “coupled to” the processor.

[0220] To provide for interaction with a user, embodiments of the subject matter described in this specification can be implemented on, or configured to communicate with, a computer having a display device, e.g., a LCD (liquid crystal display) monitor, for displaying information to the user, and an input device by which the user can provide input to the computer, e.g., a keyboard and a pointing device, e.g., a mouse, a trackball or touchpad. Other kinds of devices can be used to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input. In addition, a computer can interact with a user by sending documents to and receiving documents from a device that is used by the user; for example, by sending web pages to a web browser on a user’s device in response to requests received from the web browser, or by interacting with an app running on a user device, e.g., a smartphone or electronic tablet. Also, a computer can interact with a user by sending text messages or other forms of message to a personal device, e.g., a smartphone that is running a messaging application, and receiving responsive messages from the user in return.

[0221] This specification uses the term “configured to” in connection with systems, apparatus, and computer program components. For a system of one or more computers to be configured to perform particular operations or actions means that the system has installed on it software, firmware, hardware, or a combination of them that in operation cause the system to perform the operations or actions. For one or more computer programs to be configured to perform particular operations or actions means that the one or more programs include instructions that, when executed by data processing apparatus, cause the apparatus to perform the operations or actions. For special-purpose

logic circuitry to be configured to perform particular operations or actions means that the circuitry has electronic logic that performs the operations or actions.

[0222] While this specification contains many specific embodiment details, these should not be construed as limitations on the scope of what is being claimed, which is defined by the claims themselves, but rather as descriptions of features that may be specific to particular embodiments. Certain features that are described in this specification in the context of separate embodiments can also be realized in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be realized in multiple embodiments separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially be claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claim may be directed to a subcombination or variation of a subcombination.

[0223] Similarly, while operations are depicted in the drawings and recited in the claims in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system modules and components in the embodiments described above should not be understood as requiring such separation in all embodiments, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

[0224] Particular embodiments of the subject matter have been described. Other embodiments are within the scope of the following claims. For example, the actions recited in the claims can be performed in a different order and still achieve desirable results. As one example, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In some cases, multitasking and parallel processing may be advantageous.

What is claimed is:

CLAIMS

1. A computer-implemented method for performing counterfactual regret minimization (CRF) for strategy searching in strategic interaction between two or more parties, the method comprising:

storing a plurality of regret samples in a first data store, wherein each of the plurality of regret samples comprises a state of a party and a regret value of a possible action in the state of the party, wherein the plurality of regret samples are obtained in two or more iterations of a counterfactual regret minimization (CRF) algorithm in strategy searching in strategic interaction between the party and at least another party;

storing a plurality of strategy samples in a second data store, wherein each of the plurality of strategy samples comprises a state of the party and a strategy value of a possible action in the state of the party;

updating parameters of a first neural network for predicting a regret value of a possible action in a state of the party based on the plurality of regret samples in the first data store; and

updating parameters of a second neural network for predicting a strategy value of a possible action in a state of the party based on the plurality of strategy samples in the second data store.

2. The method of claim 1, further comprising:

identifying a first state of the party and a first possible action in the first state of the party;

predicting a first regret value of the first possible action in the first state of the party using the parameters of the first neural network; and

predicting a first strategy value of the first possible action in the first state of the party using the parameters of the second neural network.

3. The method of claim 1, further comprising:

obtaining a new regret sample; and

storing the new regret sample into the first data store according to a reservoir sampling algorithm.

4. The method of claim 3, wherein storing the new regret sample into the first data store according to a reservoir sampling algorithm comprises:

determining whether the first data store is full; and

in response to determining that the first data store is full, replacing one of the plurality of regret samples in the first data store with the new regret sample.

5. The method of claim 1, wherein the CRF algorithm comprises a robust sampling CRF algorithm.

6. The method of claim 1, wherein the strategy value of a possible action in the state of the party comprises a numerator of an average strategy.

7. The method of claim 1, wherein the regret value of a possible action in the state of the party comprises a counterfactual regret value of the possible action in the state of the party calculated based on a counterfactual value of the possible action in the state of the party.

8. The method of claim 1, further comprising:
in each of the two or more iterations of a counterfactual regret minimization (CRF) algorithm in strategy searching in strategic interaction between the party and the at least another party,
sampling a possible action out of a plurality of possible actions in a second state of a party according to a sampling scheme;
calculating a counterfactual value of the possible action in the second state of the party;
calculating a regret value of the possible action in the second state of the party based on the counterfactual value of the possible action in the second state of the party;
calculating an updated strategy of the possible action in the second state of the party based the regret value of the possible action in the second state of the party according to a regret matching algorithm; and
calculating a strategy value of the possible action in the second state of the party based on the updated strategy of the possible action in the second state of the party.

9. The method of claim 1, further comprising:
initializing the parameters of the first neural network based on parameters of the first neural network in a previous iteration; and
initializing the parameters of the second neural network based on parameters of the second neural network in a previous iteration.

10. An apparatus for performing counterfactual regret minimization (CRF) for strategy searching in strategic interaction between two or more parties, comprising:

- a first storing module for storing a plurality of regret samples in a first data store, wherein each of the plurality of regret samples comprises a state of a party and a regret value of a possible action in the state of the party, wherein the plurality of regret samples are obtained in two or more iterations of a counterfactual regret minimization (CRF) algorithm in strategy searching in strategic interaction between the party and at least another party;

- a second storing module for storing a plurality of strategy samples in a second data store, wherein each of the plurality of strategy samples comprises a state of the party and a strategy value of a possible action in the state of the party;

- a first updating module for updating parameters of a first neural network for predicting a regret value of a possible action in a state of the party based on the plurality of regret samples in the first data store; and

- a second updating module for updating parameters of a second neural network for predicting a strategy value of a possible action in a state of the party based on the plurality of strategy samples in the second data store.

11. The apparatus of claim 10, further comprising:

- an identifying module for identifying a first state of the party and a first possible action in the first state of the party;

- a first predicting module for predicting a first regret value of the first possible action in the first state of the party using the parameters of the first neural network; and

- a second predicting module for predicting a first strategy value of the first possible action in the first state of the party using the parameters of the second neural network.

12. The apparatus of claim 10, wherein the first storing module is capable of obtaining a new regret sample; and storing the new regret sample into the first data store according to a reservoir sampling algorithm.

13. The apparatus of claim 12, wherein storing the new regret sample into the first data store according to a reservoir sampling algorithm comprises:

- determining whether the first data store is full; and

- in response to determining that the first data store is full, replacing one of the plurality of regret samples in the first data store with the new regret sample.

14. The apparatus of claim 10, wherein the CRF algorithm comprises a robust sampling CRF algorithm.

15. The apparatus of claim 10, wherein the strategy value of a possible action in the state of the party comprises a numerator of an average strategy.

16. The apparatus of claim 10, wherein the regret value of a possible action in the state of the party comprises a counterfactual regret value of the possible action in the state of the party calculated based on a counterfactual value of the possible action in the state of the party.

17. The apparatus of claim 10, further comprising:
in each of the two or more iterations of a counterfactual regret minimization (CRF) algorithm in strategy searching in strategic interaction between the party and the at least another party,
a sampling module for sampling a possible action out of a plurality of possible actions in a second state of a party according to a sampling scheme;
a first calculating module for calculating a counterfactual value of the possible action in the second state of the party;
a second calculating module for calculating a regret value of the possible action in the second state of the party based on the counterfactual value of the possible action in the second state of the party;
a third calculating module for calculating an updated strategy of the possible action in the second state of the party based the regret value of the possible action in the second state of the party according to a regret matching algorithm; and
a fourth calculating module for calculating a strategy value of the possible action in the second state of the party based on the updated strategy of the possible action in the second state of the party.

18. The apparatus of claim 10, further comprising:
a first initializing module for initializing the parameters of the first neural network based on parameters of the first neural network in a previous iteration; and
a second initializing module for initializing the parameters of the second neural network based on parameters of the second neural network in a previous iteration.

19. A system, comprising:

one or more processors; and

one or more computer-readable memories coupled to the one or more processors and having instructions stored thereon which are executable by the one or more processors to perform the method of any of claims 1-9.

20. A computer-implemented method for performing counterfactual regret minimization (CRF) for solving an imperfect information game (IIG) that involves two or more parties, the method comprising:

storing a plurality of regret samples in a first data store, wherein each of the plurality of regret samples comprises a state of a party and a regret value of a possible action in the state of the party, wherein the plurality of regret samples are obtained in two or more iterations of a counterfactual regret minimization (CRF) algorithm in strategy searching in strategic interaction, wherein the IIG represents a collaborative product-service recommendation service that involves at least the party and a second party, the party having limited access to information of the second party, wherein the state of the party comprises a history of information provided by the second party, and wherein the possible action of the party comprises a possible action in response to the history of information provided by the second party for providing product-service recommendations to customers;

storing a plurality of strategy samples in a second data store, wherein each of the plurality of strategy samples comprises a state of the party and a strategy value of a possible action in the state of the party;

updating parameters of a first neural network for predicting a regret value of a possible action in a state of the party based on the plurality of regret samples in the first data store; and

updating parameters of a second neural network for predicting a strategy value of a possible action in a state of the party based on the plurality of strategy samples in the second data store.