Artificial Swarm Intelligence vs Vegas Betting Markets

Louis Rosenberg
Unanimous AI
San Francisco, CA, USA
Louis@Unanimous.ai

Gregg Willcox
Unanimous AI
San Francisco, CA, USA
Gregg@Unanimous.ai

***Abstract*— In the natural world, Swarm Intelligence (SI) is a commonly occurring process in which biological groups amplify their collective intelligence by forming closed-loop systems. It is well known in schools of fish, flocks of bird, and swarms of bees. In recent years, new AI technologies have enabled networked human groups to form systems modeled after natural swarms. Known as Artificial Swarm Intelligence (ASI), the technique has been shown to amplify the effective intelligence of human groups. This study compares the predictive ability of ASI systems against large betting markets when forecasting sporting events. Groups of average sports fans were tasked with predicting the outcome of 200 hockey games (10 games per week for 20 weeks) in the NHL. The expected win rate for Vegas favorites was 62% across the 200 games based on the published odds. The ASI system achieved a win rate of 85%. The probability that the system outperformed Vegas by chance was extremely low (p = 0.0057), indicating a significant result. In addition, researchers compared the winnings from two betting models – one that wagered weekly on the Vegas favorite, and one that wagered weekly on the ASI favorite. At the end of 20 weeks, the Vegas model generated a 41% financial loss, while the ASI model generated a 170% financial gain.**

Keywords*— *Swarm Intelligence, Artificial Swarm Intelligence, Collective Intelligence, Human Swarming, Artificial Intelligence.

I. BACKGROUND

Prior studies on Artificial Swarm Intelligence (ASI) have shown that by forming real-time "human swarms," networked human groups can significantly amplify their accuracy in a wide variety of forecasting tasks [1-6], outperforming traditional "Wisdom of Crowd" methods [3]. For example, a 2015 study assessed the ability of human swarms to forecast the outcome of college football games. An ASI swarm comprised of 75 amateur sports fans, connected by AI algorithms, was tasked with predicting 10 bowl games. As individual forecasters, the participants averaged 50% accuracy when predicting outcomes against the Vegas spread. When forecasting as a real-time ASI system, those same participants achieved 70% accuracy against the Vegas spread [2]. Similar increases have been demonstrated in other studies, including a 5-week study that tasked human participants, connected as an ASI system, with predicting 50 consecutive soccer matches in the English Premier League. Results showed a 31% increase in accuracy when participants were connected in ASI swarms [4]. The ASI system also outperformed the BBC's machine-model known as "SAM" over the same 50 games. [13].

While prior studies have documented the ability of artificial swarms to outperform individuals and outperform traditional Wisdom of Crowd methods across a range of forecasting tasks, no formal study has compared the predictive ability of artificial swarms against largescale markets. To address this need, a study was run to compare human swarms to Vegas betting markets, assessing the accuracy rates and the financial returns across a large set of predictions. Specifically, this study required human participants to forecast the outcome of 200 games in the National Hockey League (NHL), structured as 10 games per week for 20 consecutive weeks.

II. SWARMS VS CROWDS

When comparing the accuracy of real-time swarms against traditional crowd-based methods, it's worth reviewing the structural differences between them. The prime differentiator between "crowds" and "swarms" is that in crowd-based methods, human participants provide input in isolation for aggregation in external statistical models, while in swarm-based methods, human participants "think together" in real-time, their interactions governed by intelligence algorithms. This means that swarms are closed-loop systems in which participants act, react, and interact with each other, converging on optimized solutions in synchrony. The swarming process is generally modeled after biological systems such as schools of fish, flocks of birds, and swarms of bees. The present study uses Swarm AI technology from Unanimous AI Inc, which is modeled largely on the decision-making processes of honeybee swarms [4].

As background, the decision-making processes that govern the behavior of honeybee swarms have been studied since the 1950s and have revealed themselves to be very similar to the decision-making processes in neurological brains [7-9]. Both brains and swarms employ large populations of simple excitable units (i.e., bees and neurons) that operate in parallel to integrate noisy evidence, weigh competing alternatives, and converge on decisions in synchrony. In both, outcomes are reached through real-time competition among sub-populations of excitable units. When the support generated by one sub-population exceeds a threshold level, that alternative is chosen. In honeybees, this enables the group to converge on optimal decisions, picking the best solution to complex problems 80% of the time [11,12].

978-1-5386-6712-5/18/$31.00 ©2018 IEEE

III. ENABLING "HUMAN SWARMS"

Unlike birds and bees and fish, humans have not evolved the natural ability to form closed-loop systems that enables real-time swarming. We lack the subtle connections that other organisms use to establish high speed feedback-loops among members. Schooling fish detect vibrations in the water around them. Flocking birds detect subtle motions propagating through the population. Swarming bees use complex body vibrations called a "waggle dance." To enable real-time swarming among groups of networked humans, specialized software is required to close the loop among all members.

To address this need, a software platform (swarm.ai) was developed to enable networked human populations to form real-time swarms by connecting from anywhere in the world [1]. Modeled on the decision-making process of honeybee swarms, the cloud-based swarm.ai system enables groups of users to work in parallel to (a) integrate noisy evidence, (b) weigh competing alternatives, and (c) converge on decisions in synchrony, while also allowing all participants to perceive and react to the changing system in real-time, thereby closing a feedback loop around the full population of participants.

As shown in Figure 3 below, the human participants of ASI systems answer questions by moving a graphical puck to select among a set of alternatives. Each participant provides input by manipulating a graphical magnet with a mouse, touchscreen, or other input device. By positioning their magnet with respect to the moving puck, real-time participants express their personal intent, impacting the system as a whole. The input from each user is not a discrete vote, but a continuous stream of vectors that varies freely over time. Because all members of the swarming population can adjust their intent fluidly in real-time, the ASI swarm explores the decision-space, not based on the input of any individual, but based on the emergent dynamics of the full system. This enables complex deliberations across all members at once, empowering the group to collectively explore all the options and converge upon the one solution that best represents their combined insights.

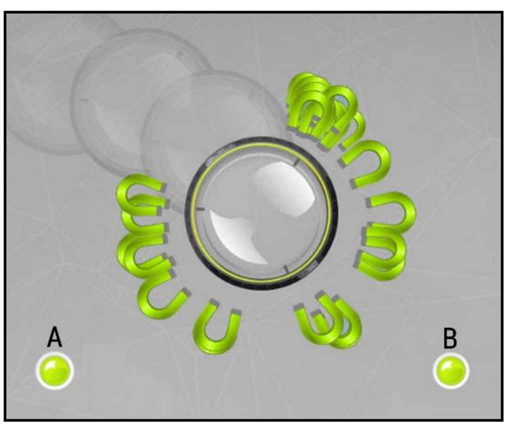

Fig. 3. A human swarm choosing between options

It is important to note that participants do not only vary the direction of their intent, but also modulate the magnitude of their intent by adjusting the distance between their magnet and the puck. Because the puck is in continuous motion across the decision-space, users need to continually move their magnet so that it stays close to the puck's outer rim. This is significant, for it requires participants to be engaged continuously throughout the decision process, evaluating and re-evaluating their intent. If they stop adjusting their magnet with respect to the changing position of the puck, the distance grows and their applied sentiment wanes.

Thus, like bees vibrating their bodies to express sentiment in a biological swarm, or neurons firing activation signals to express conviction levels within a biological neural-network, the participants in an artificial swarm must continuously update and express their changing preferences during the decision process, or lose their influence over the collective outcome. In addition, intelligence algorithms monitor the behaviors of all swarm members in real-time, inferring their implied conviction based upon their relative motions over time. This reveals a range of behavioral characteristics within the swarm population and weights their contributions accordingly, from entrenched participants to flexible participants to fickle participants.

IV. PREDICTION STUDY

To assess the ability of human swarms to outperform Vegas betting markets, a formal study was conducted over a 20-week period using groups of randomly selected human subjects from a pool of self-reported sports enthusiasts. Each weekly group consisted of 25 to 36 participants, all of whom logged in remotely to the cloud-based *swarm.ai* system. Human subjects were paid $3.00 for their participation in each weekly session, which required them to make predictions of the outcome of all ten hockey games being played that night, participating both as (a) individuals reporting on a standard online survey, and (b) as part of a real-time ASI system.

For each hockey game, participants were tasked with forecasting the winner and the margin of victory, expressed as either (a) the team win by 1 goal, or (b) the team win by 2 or more goals. The margins were chosen to match common Vegas gambling spreads. Figure 4 below shows a snapshot of a human swarm comprised of 31 participants in the process of predicting a match between Toronto and Calgary.

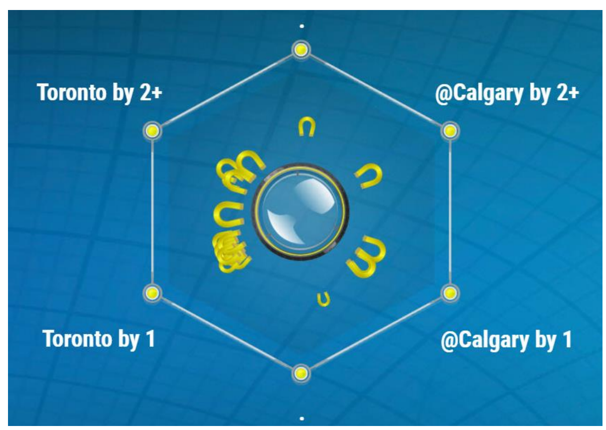

Fig. 4. Human Swarm in the process of forecasting an NHL game

As shown in Figure 4, each real-time swarm is tasked with selecting from among four outcome options, indicating which team will win and which margin is most likely. Again, the pparticipants do not cast discrete votes but express their intent continuously over time, converging together as a system. The image shown in Figure 4 is a snapshot of the system as it moves across the decision-space and converges upon an answer, a process that generally takes between 10 and 60 seconds.

In addition to forecasting each individual game, participants were asked to identify which of the weekly predictions is the most likely to be a correct assessment. In other words, which of the teams forecast to win their games should be deemed the "pick of the week" by virtue of being the most likely to win their game. Figure 5 below shown an example ASI system in the process of identifying the pick of the week.

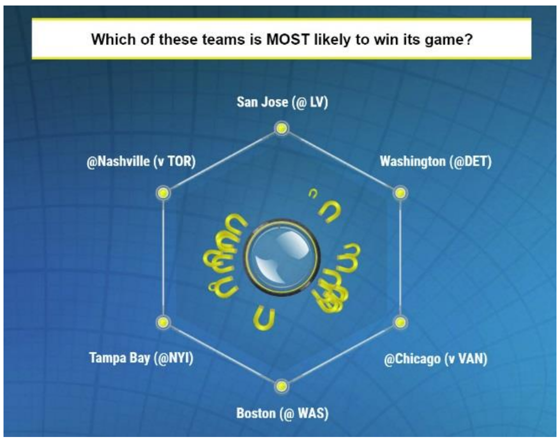

Fig. 5. Human Swarm in process of identifying "Pick of the Week"

V. Wagering Protocol

By collecting predictions for each of the 10 weekly games as well as a top "pick of the week", forecasting data was collected across all 20 weeks for accuracy comparison against Vegas betting markets. To enable ROI comparisons against betting markets, two standardized betting models were tracked across the 20-week period. In both models, an initial simulated betting pool of $100 was created as the starting point for ROI computations, the pools tracked over the 20-week period.

In "Wagering Model A," a simple heuristic was defined which allocated weekly bets equal to 15% of the current betting pool, dividing it equally across all ten weekly forecasts made by the ASI system. In "Wagering Model B," a similar heuristic was defined which also allocated 15% of the current betting pool for use in weekly bets, but placed the entire 15% upon the one game identified as "pick of the week". Both pots were tracked over the 20-week period, using actual Vegas payouts to compute returns. Vegas odds used in this study were captured from www.sportsbook.ag, a popular online betting market.

VI. Results

Across the set of 200 games forecast by the ASI system, an accuracy rate of 61% was achieved. This compares favorably to the expected accuracy of 55% based on Vegas odds (p=0.0665). Of course, the more important skill in forecasting sporting events is identifying which games can be predicted with high confidence as compared to those games which are too close to call. This skill is reflected in the "pick of the week" generated by the ASI system. Across the 20 weeks, the system achieved 85% accuracy in correctly predicting the winner of the "pick of the week" game. This compares very favorably to the expected accuracy of 62% based on Vegas odds.

Figure 6 below shows the distribution of Vegas Odds for the twenty selected "pick of the week" games. As described above, the swarm-based system had a win rate of 85% across these same games. This is a significant improvement, equivalent to reducing the error in Vegas Odds by 61%. The probability that the swarm outperformed Vegas Odds by chance was extremely low (p = 0.0057), indicating a highly significant result.

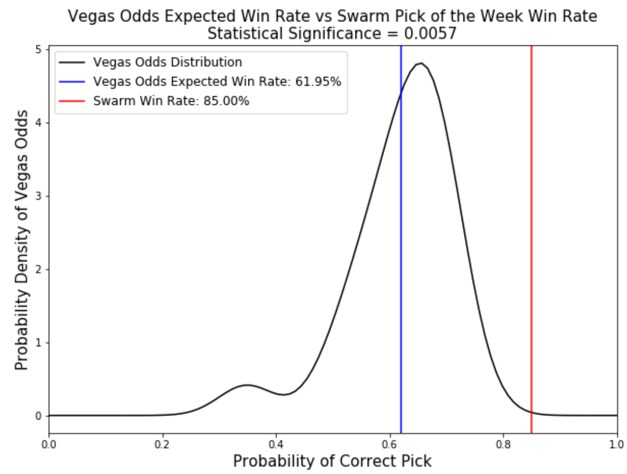

Fig 6. Summary of results across 20 weeks of NHL predictions

In addition, a betting simulation was run for each prediction set in which 15% of the current bankroll was bet on each prediction in each week. The performance of this model when betting against Vegas (and including the Bookie's cut) is seen below in Figure 7. Starting with $100 and investing each week according to this strategy, the Pick of the Week strategy results in a gain of $270.20, equivalent to a 20-week ROI of 170%, and a week-over-week average ROI of 5.09%. For comparison, betting on all of the swarm's picks evenly (for a total of 15% of the bankroll) results in $121.82, or a 20-week ROI of 21.8%, indicating that the swarm is selecting better than randomly among its picks.

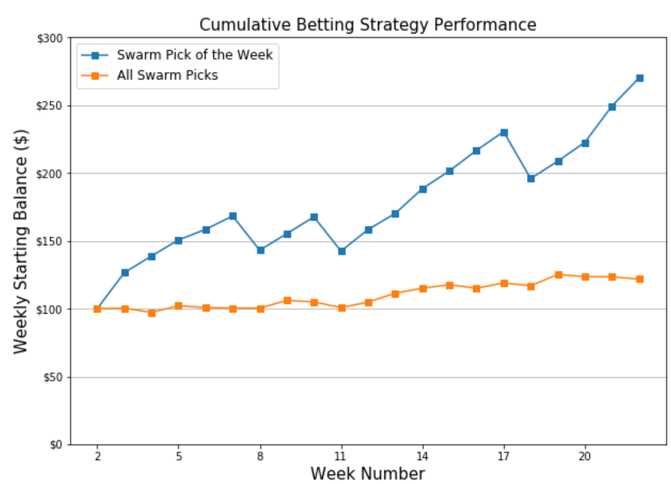

Fig 7. Cumulative Betting Performance across 20 weeks

While the above results are impressive, especially the 170% ROI over 20 weeks, we can gain additional insight into the significance of this outcome by comparing against additional baselines. For example, we can (a) compare these results to randomly placed bets across all games played as a means of assessing if the swarm bets across all games are as significant as they appear, and (b) compare these results to bets placed on the Vegas favorite each week as a means of assessing if betting on the swarm's top picks each week is as impressive as it seems.

These baselines are shown in Figure 8 as the green line and red line respectively. Looking first at random betting across all games, the net outcome across 20 weeks was $72.39, which equates to 28% loss over the test period. This is significantly worse than the $122 (22% gain) achieved by betting on all swarm-based forecasts. Even more surprising, betting on the Vegas favorites each week resulted in a net outcome of $59, which equates to a 41% loss over the 20-week test period. This is significantly worse than the $270 (170% gain) achieved by betting on the swarm's top picks.

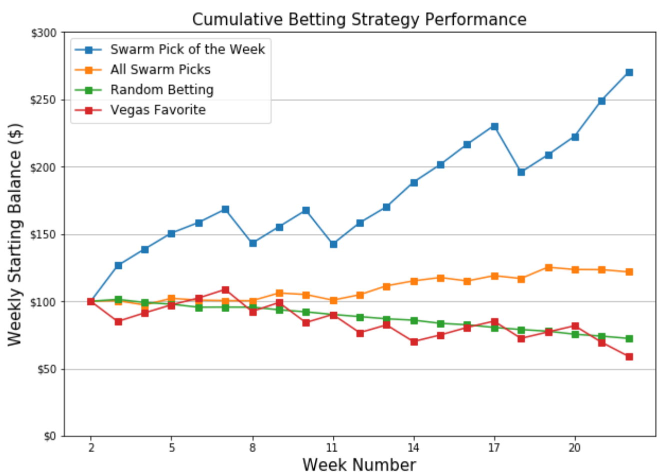

Fig 8. Swarm Performance vs Baseline Performance across 20 weeks

VII. CONCLUSIONS

Can real-time "human swarms" outperform the predictive abilities of largescale betting markets? The results of this study suggest this may be the case. As shown across a set of 200 NHL games during the 2017-2018 hockey season, an ASI system comprised of 25 to 36 average sports fans, connected by intelligence algorithms, significantly out-performed Vegas in predictive accuracy. The results were strongest when the ASI system was tasked with identifying a "pick of the week" as the most likely game to achieve the predicted outcome. Across the 20 weeks, the system achieved 85% accuracy when predicting the "pick of the week", which compares very favorably to the expected accuracy of 62% based on Vegas odds. The probability that the system outperformed Vegas by chance was extremely low (p = 0.0057), indicating a highly significant result.

In addition, when using the "pick of the week" as part of an automated wagering heuristic, a simulated betting pool that began at $100 at the start of the experiment, increased to $270 over the 20-week period based on the swarm-based predictions. This corresponds to an impressive return on investment (ROI) of 170% across the testing period. Additional work is being conducted to optimize this heuristic, as there appears to be room for improvement when generating Vegas wagers based on a swarm-based predictive intelligence.

ACKNOWLEDGMENT

We greatly appreciate the efforts of Chris Hornbostel and David Baltaxe for coordinating human participants in weekly sessions. We also acknowledge Unanimous AI for providing access to the **swarm.ai** platform.

REFERENCES

[1] Rosenberg, L.B., "Human Swarms, a real-time method for collective intelligence." Proceedings of the European Conference on Artificial Life 2015, pp. 658-659

[2] Rosenberg, Louis. "Artificial Swarm Intelligence vs Human Experts," Neural Networks (IJCNN), 2016 International Joint Conference on. IEEE. J. Clerk Maxwell, A Treatise on Electricity and Magnetism, 3rd ed., vol. 2. Oxford: Clarendon, 1892, pp.68–73.

[3] Rosenberg, Louis. Baltaxe, David and Pescetelli, Nicollo. "Crowds vs Swarms, a Comparison of Intelligence," IEEE 2016 Swarm/Human Blended Intelligence (SHBI), Cleveland, OH, 2016, pp. 1-4.

[4] Baltaxe, David, Rosenberg, Louis and N. Pescetelli, "Amplifying Prediction Accuracy using Human Swarms", Collective Intelligence 2017. New York, NY ; 2017.

[5] Rosenberg, L, Pescetelli, N, & Willcox, G. "Artificial Swarm Intelligence amplifies accuracy when predicting financial markets," 2017 IEEE 8th Annual Ubiquitous Computing, Electronics and Mobile Communication Conference (UEMCON), New York City, NY, 2017, pp. 58-62..

[6] Rosenberg, L. and Willcox, G. "Artificial Swarms find Social Optima" 2018 IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA), Boston, MA, 2018, pp. 174-178.

[7] Seeley T.D, Buhrman S.C 2001 "Nest-site selection in honey bees: how well do swarms implement the 'best-of-N' decision rule?" Behav. Ecol. Sociobiol. 49, 416–427

[8] Marshall, James. Bogacz, Rafal. Dornhaus, Anna. Planqué, Robert. Kovacs, Tim. Franks, Nigel. "On optimal decision-making in brains and social insect colonies." Soc. Interface 2009.

[9] Seeley, Thomas D., et al. "Stop signals provide cross inhibition in collective decision-making by honeybee swarms." Science 335.6064 (2012): 108-111.

[10] Usher, M. McClelland J.L 2001 "The time course of perceptual choice: the leaky, competing accumulator model." Psychol. Rev. 108, 550–592

[11] Seeley, Thomas D. Honeybee Democracy. Princeton Univ. Press, 2010.

[12] Seeley, Thomas D., Visscher, P. Kirk. "Choosing a home: How the scouts in a honey bee swarm perceive the completion of their group decision making." Behavioural Ecology and Sociobiology 54 (5) 511-520.

[13] McHale, Ian. Sports Analytics Machine (SAM) as reported by BBC: http://blogs.salford.ac.uk/business-school/sports-analytics-machine/