Human Swarms amplify accuracy in Honesty Detection

LOUIS ROSENBERG, Unanimous A.I.

1. INTRODUCTION

It goes back to the birds and the bees. Fish too. And ants. Even slime-molds. It goes back to all creatures that amplify their intelligence collectively by forming flocks, schools, colonies, or swarms. Across countless species, nature proves that social creatures, when working together in closed-loop systems, can outperform the vast majority of individual members when solving problems and making decisions. Biologists call this *Swarm Intelligence* (SI) and it has inspired a new category of A.I. that enables similar amplification effects in human domains. The technology is called *Artificial Swarm Intelligence* (ASI) and it works by connecting groups of people into closed-loop systems online, governed by A.I. algorithms. These "human swarms" can answer questions and solve problems together in synchrony [Rosenberg 2015, Eberhart 2015].

Recent studies have shown that human swarms can significantly amplify the predictive ability of groups, outperforming individual members as well as outperforming traditional collective intelligence techniques such as polls, votes, and markets. In one recent study conducted by researchers at Unanimous A.I. and Oxford University, human subjects were tasked with predicting a set of 20 official Las Vegas wagers known as "proposition bets" on Super Bowl 50. Traditional crowd-sourced predictions were pitted against the predictions made by a real-time swarm. The crowd was composed of 467 football fans who provided their predictions by online survey. The swarm was composed of 29 football fans who provided their predictions together as a closed-loop system. Although the crowd was 16 times larger than the swarm, it was far less accurate, achieving only 47% correct predictions and generating a 9% gambling loss. The swarm achieved 68% correct predictions and generated a 36% gambling win [Rosenberg, Baltaxe, and Pescetelli 2016].

While many other studies have produced similar results, demonstrating the ability of human swarms to amplify the predictive intelligence of human groups, no prior work has looked at the social intelligence of real-time swarms. To address this, the present study tested the ability of human swarms to assess honesty in human faces. Specifically, the study compared the ability of individuals with the ability of swarms when tasked with identifying if a smiling person was producing a *"real smile"* authentically (i.e. evoked in response to a joyful stimuli), or producing a *"fake smile"* deceptively (i.e. evoked artificially on demand).

1.1 Assessing Honesty by Judging Smiles

In 1969, Ekman and Friesen published the first critical research into the correlation between facial cues and human deception. They defined "leakage cues" as involuntary expressions that reveal if a person's true feelings don't match what they're consciously attempting to convey [Ekman and Friesen, 1969]. Smiles have been identified as a significant leakage cue that can be used to determine if a person is being honest or deceitful, especially if they are faking emotions [Ekman, Friesen, and Davidson 1990].

While the difference between genuine smiles and fake smiles can be measured scientifically, most people are not very good at telling the difference. The facial cues are often subtle and therefore easily missed by the untrained eye. [Ekman, Friesen, and O'Sullvan, 1988. DePaulo, 2003]. In a recent study, a group of 217 subjects were asked to view a set of 20 videos of individuals smiling. The videos represented a mix of genuine smiles (produced by enjoyment) and fake smiles (produced on demand). Across the 217 subjects, the average person's assessment was incorrect 32% of the time [Kajdasz, 2014].

The question thus remains - can Artificial Swarm Intelligence be used to reduce the error rates in facial assessment tasks? If so, it may suggest a pathway to amplifying our ability to assess human honesty.

2. SWARMS AS INTELLIGENT SYSTEMS

The most deeply studied swarms in nature are honeybee swarms, which are known to find optimized solutions to complex multi-variable problems. In fact, the decision-making processes in honeybee swarms have been shown to be remarkably similar to the decision-making processes in neurological brains [Seeley, 2010; Seeley et al., 2012]. Both employ large populations of simple excitable units (i.e., bees and neurons) that work in parallel to integrate noisy evidence, weigh competing alternatives, and converge on decisions in synchrony. In both, outcomes are arrived at through a real-time competition among sub-populations of excitable units, each sub-population vying for one of a plurality of alternate solutions. When one sub-population exceeds a threshold level of support, the corresponding alternative is chosen. The threshold in both brains and swarms is not unanimous support of the population, or even a simple majority, but a sufficient quorum of excitation. In honeybees, this helps to avoid deadlocks and results in optimal decisions over 80% of the time [Seeley, 2010; Seeley et al., 2012]. Thus by working together as a real-time unified system, the bee colony amplifies its intelligence beyond the capacity of individual members. It is this amplification of intelligence that human swarming aims to enable among groups of networked people.

2.1 Enabling Human Swarms

To amplify the intelligence of online human groups, specialized technologies are required to close the loop among members. To address this need, an online platform called UNU was developed to allow distributed groups of users to login from anywhere in the world and participate in closed loop swarming processes. Modeled after the decision-making of honeybee swarms, the UNU platform allows groups of independent actors to work in parallel to (a) integrate noisy evidence, (b) weigh competing alternatives, and (c) converge on final 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.

As shown in Figure 1 below, participants in the UNU platform answer questions by collectively moving a graphical puck to select among a set of alternatives. The puck is modeled as a physical system with a defined mass, damping and friction. Each participant provides input by manipulating a graphical magnet with a mouse or touchscreen. By positioning their magnet, users impart their personal intent as a force vector on the puck. The input from each user is not a discrete vote, but a stream of vectors that varies freely over time. Because the full populations of users can adjust their intent at every time-step (200 ms), the puck moves, not based on the input of any individual, but based on the dynamics of the full system. This enables real-time physical negotiation among all members, empowering the group to collectively explore the decision-space and converge on the most agreeable answer in synchrony.

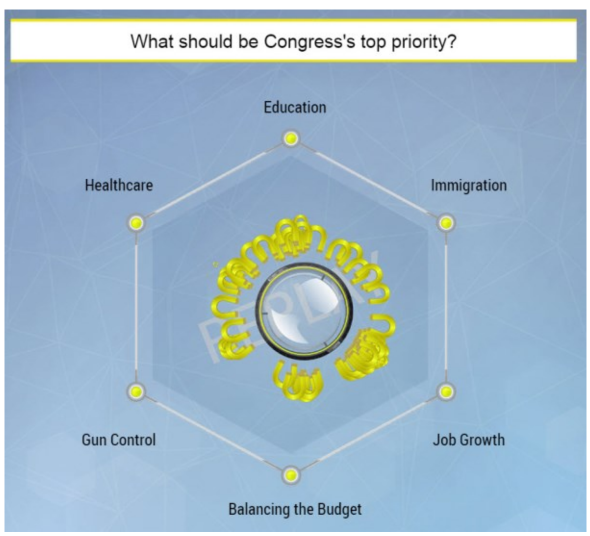

Fig. 1. Shown is a snapshot of a human swarm in the process of answering a question in real-time.

It's important to note that participants can only see *their own magnet* during the decision process**,** and not the magnets of others members of the swarm. Thus, although they can view the puck's motion in real time, which represents the emerging will of the full system, they are not influenced by the specific breakdown of support across the options. This greatly reduces social influence biasing. For example, if the puck slows due to an emerging deadlock, the participants must evaluate their own willingness to shift support to alternate solutions without knowing the specific distribution of support that caused the deadlock.

It's also important to note that users don't just vary the direction of their input, but also the magnitude by adjusting the distance between the magnet and the puck. Because the puck is in motion, to apply full force users need to continually move their magnet so that it stays close to the puck's rim. This is significant, for it requires all users to be engaged during the decision process. If they stop adjusting their magnet to the changing position of puck, the distance grows and their applied force wanes. Thus, like bees vibrating their bodies or neurons firing activation signals, the users in an artificial swarm must continuously express their changing preferences during the decision process or lose their influence over the collective outcome.

3. HUMAN SWARMS AND SMILE ASSESSMENT

To address the research question of whether an Artificial Swarm Intelligence comprised of distributed online participants can assess the authenticity of smiles with increased accuracy as compared to individual human assessors, a formal research study was conducted. Five groups of test subjects were randomly recruited of mixed age and gender. Each group had 30 to 35 adult members, all of whom logged in remotely using a personal computer and none of whom expressed any special training in smile assessment.

The participants of each of test group were tasked with viewing and assessing a set of 20 smile videos. Each smile video presented a 3 second clip of a person forming a smile in response to stimuli. The videos were sourced from existing research by smile expert Paul Eckman [Eckman, 2014]. Figure 2 below shows snapshots from two of the smile videos. The snapshot on the left depicts a joyful smile that was generated in response to an authentic stimulus. The snapshot on the right depicts a fake smile that was produced on demand, and not in response to a joyful stimulus.

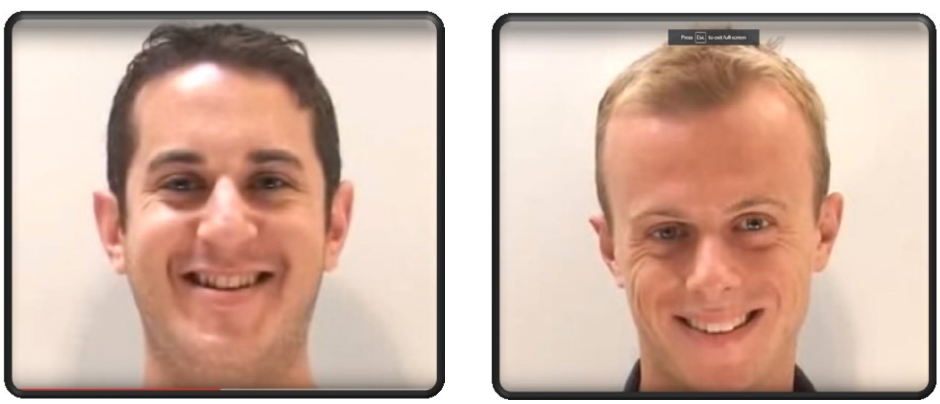

Authentic "Joyful" Smile   Inauthentic "Fake" smile

Fig. 2. Shown are snapshot from two sample smile videos used in the assessment study.

Upon viewing each of the videos, every test subject was required to express their personal assessment as to whether the video depicted a genuine "joyful smile" or a deceitful "fake smile." This was performed twice for each individual – once by working alone and reporting their assessment on a standard online survey, and once by working together with other members of their group as an Artificial Swarm Intelligence.

Figure 3 below shows a snapshot of one of the three groups in the process of assessing one of the twenty videos by working as a swarm. As shown, all members of the 35 person group are working together to move the glass puck by individually positioning and repositioning their graphical magnets in synchrony. In this way, the group explores the decision-space and converges on a single unified assessment.

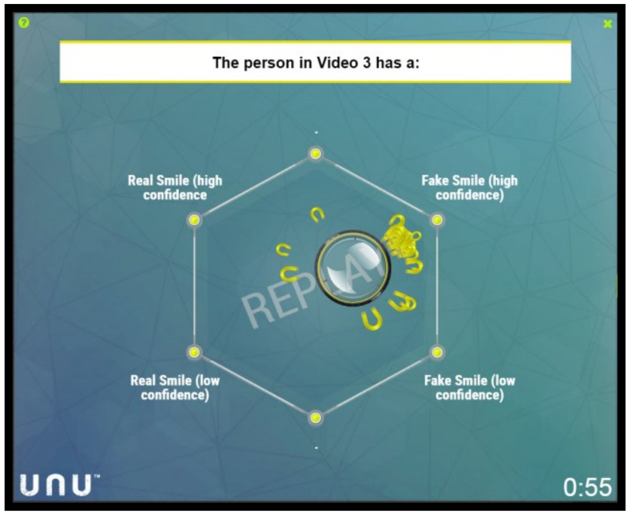

Fig. 3. Shown is a snapshot of a human swarm of 35 participants in the process of assessing a smile video in real-time.

It should be noted that each swarm was limited to only 60 seconds for viewing and assessing a single video, with most assessments being executed in under 20 seconds. It should also be noted that videos were only played once – subjects were not allowed to repeatedly review a video as they assessed the authenticity of the smile. It should also be noted that although Figure 3 shows all of the magnets of all of the members of the swarming group, the participants could only see their own magnet during the decision process.

3.1 Results

The table and graph in Figure 4 shows the results for all five trials of the smile assessment test. Each trial was comprised of 20 video assessments executed by a unique group of test subjects. Across the five trials, the average number of errors across individual assessors was 6.6 incorrect assessments. When those same individuals worked together as a group-wise swarm, the average error rate was reduced to 3.6 incorrect assessments. This is a statistically significant improvement (P-value=0.1110) and corresponds with an average error reduction from swarming of **46%** across trials (± 16%). In other words, by working together as real-time swarms, each group of participants were able to assess the authenticity of smiles with 46% fewer errors, on average, as compared to individuals working alone.

	Group Size	INDIVIDUALS Avg # of Errors	SWARM # of Errors	Error Reduction
Trial 1	33	5.73	3	48%
Trial 2	35	7.03	4	43%
Trial 3	35	6.64	5	25%
Trial 4	30	6.40	2	69%
Trial 5	35	7.31	4	45%
Average	**34**	**6.62**	**3.60**	**46%**
		σ = 0.61	σ = 1.14	σ = 16%

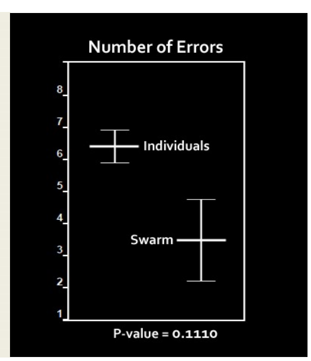

Fig 4. Results across five trials of the smile assessment test, each trial comprised of 20 smile video assessments.

Although this study aims to compare the performance of human swarms to individual assessors, we can also compare the performance of the swarm to the traditional crowd-sourcing method of aggregating poll results across sets of independent respondents. Doing this for each of the five trials, the most popular poll result across the members of each group was used as the final smile-assessment for that group. This produced an average of **5.3** incorrect assessments across the five trials, which is still a significantly higher error rate than the **3.5** assessment errors for swarm-based responses. Specifically, swarming decreased the smile assessment errors by **33%** as compared to the traditional "Wisdom of Crowd" methodology.

3.2 Conclusions

The results above suggest that forming an Artificial Swarm Intelligence comprised of a plurality of distributed users can significantly reduce the error rate when evaluating human smiles for authenticity. This suggests that human swarms are not only useful for amplifying predictive intelligence, but is also useful for amplifying social intelligence. In addition, because smile authenticity is a direct indicator of human honesty vs. deceit, this opens a range of possible applications of ASI technology from intelligence screening to jury selection. Future research is required to further explore the ability of human swarms to identify deceit, not just in facial expression but in verbal statements. In addition, further research is required to determine the ideal number of participants for a deceit-detecting swarm.

ACKNOWLEDGEMENTS

Many thanks to David Baltaxe and Chris Hornbostel of Unanimous A.I. for their significant efforts in supporting the data collection for this research study. Thanks also for the use of the unu.ai platform.

REFERENCES

Eberhart, Russell, Daniel Palmer, and Marc Kirschenbaum. "Beyond computational intelligence: blended intelligence." Swarm/Human Blended Intelligence Workshop (SHBI), 2015. IEEE, 2015.

Ekman, Paul. and Friesen, Wallace. The repertoire of non-verbal behavior: categories, origin, usage, and coding. Semiotica, 1. 49-98. 1969

Ekman, Paul., Friesen, Wallace., and O'Sullivan, Maureen. Smiles when lying. Journal of Personality and Social Psychology, 1988, Vol 54, 414-420.

Ekman, Paul., Friesen, Wallace., and Davidson, Richard. The Duchenne Smile: Emotional Expression and Brain Physiology. Journal of Personality and Social Psychology, 1990, Vol 58, No 2, 342-353.

Ekman, Paul (2014)- Spot the Fake Smile: http://www.ekmaninternational.com/paul-ekman-international-plc-home/news/can-you-spot-a-fake-smile.aspx

Kajdasz, James. A demonstration of the Benefits of Aggregation in an Analytic Task. International Journal of Intelligence and Counter Intelligence. 2014. 27: 752-763.

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

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

Rosenberg, Louis and Pescetelli, Nicollo. Amplifying Prediction Accuracy using Swarm AI, Intelligent Systems Conference (IntelliSys) 2017, London, UK. Pending.

Rosenberg, Louis. Artificial Swarm Intelligence vs Human Experts, IEEE International Joint Conference On Neural Networks (IJCNN), 2016.

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

Seeley, Thomas D. Honeybee Democracy. Princeton Univ. Press, 2010.

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. Behavioral Ecology and Sociobiology 54 (5) 511-520.