

Crowdsourcing Controller - Utilizing Reliable Agents in a Multiplayer Game

1st Kacper Kenji Lesniak

University of Copenhagen and Dynasty Studios
Copenhagen, Denmark
kkl@di.ku.dk

2nd Maria Maistro

University of Copenhagen
Copenhagen, Denmark
mm@di.ku.dk

Abstract—This paper presents a new use case for continuous crowdsourcing, where multiple players collectively control a single character in a video game. Similar approaches have already been proposed, but they suffer from certain limitations: (1) they simply consider static time frames to group real-time inputs from multiple players; (2) then they aggregate inputs with simple majority vote, i.e., each player is uniformly weighted.

We present a continuous crowdsourcing multiplayer game equipped with our Crowdsourcing Controller. The Crowdsourcing Controller addresses the above-mentioned limitations: (1) our Dynamic Input Frame approach groups incoming players' input in real-time by dynamically adjusting the frame length; (2) our Continuous Reliability System estimates players' skills by assigning them a reliability score, which is later used in a weighted majority vote to aggregate the final output command.

We evaluated our Crowdsourcing Controller offline with simulated players and online with real players. Offline and online experiments show that both components of our Crowdsourcing Controller lead to higher game scores, i.e., longer playing time. Moreover, the Crowdsourcing Controller is able to correctly estimate and update players' reliability scores.

Index Terms—Crowdsourcing, Multiplayer, Reliability system

I. INTRODUCTION

Crowdsourcing systems are fascinating and powerful tools that can be utilized in a variety of use cases [3]. Thanks to wide access to a stable and fast network connection, it is possible now to develop those systems for real-time applications, i.e., continuous crowdsourcing. In recent years, research has been dedicated to real-time and online crowdsourcing systems, with use cases such as video captioning, robot controls, audio processing, and assistive technology [8].

Another topic that emerged in recent years are video game streams, where spectators become active players by voting on certain gameplay choices or even taking direct control of the game. Recent research has found that watching a stream of a video game is more about social interaction than the game content itself [2]. This type of games, where larger crowds play and cooperate together through aggregated network input are described as Emergent Multiplayer Games (EMG) [14]. The popularity of these events shows that there is a market for new, innovative ways in which players can cooperate in game contexts. Although research has been put into how a

multiplayer voting mechanisms can work from a gameplay perspective [14], there is little insight into merging crowdsourcing approaches with multiplayer video game controllers.

In this paper, we explore continuous crowdsourcing approaches in the context of multiplayer games, in which players make collective decisions regarding the gameplay. We present a new use case for continuous crowdsourcing, where multiple players collectively control a single character in a video game. Similar approaches have already been proposed in different contexts, but they suffer from certain limitations. They consider static frames to group inputs from multiple users and they do not utilize approaches commonly implemented in crowdsourcing systems to estimate players reliabilities when aggregating inputs.

The contribution of this paper is twofold. The first part is a simple video game with an innovative mechanic that lets multiple network players control one character through an aggregated input. The second part is our Crowdsourcing Controller, which controls how to: (1) group real-time inputs coming from multiple users into frames; (2) aggregate inputs within a frame and issue a single output command. We propose two novel approaches to aggregate and process crowdsourced inputs: the Dynamic Input Frame to group inputs and the Continuous Reliability System to aggregate them. The former aims at detecting a precise moment in time when the input should be issued to the game, while the latter focuses on recognizing and utilizing reliable players during the play session. A combination of offline, AI-based experiments, as well as online experiments, with 9 real players, has been conducted to gather the data needed for the analysis and evaluation of the proposed solutions.

The paper is organized as follows: related work on crowdsourcing and online multiplayer games (§ II); the online game (§ IV); our Crowdsourcing Controller (§ IV); experimental evaluation (§ V); conclusions and future work (§ VI).

II. RELATED WORK

First, we describe crowdsourcing, its definitions and most prevalent use cases (§ II-A). Then we present related work on Emergent Multiplayer Games (EMG), a type of online video streams in which viewers have direct control over the game through voting mechanisms (§ II-B).

A. Crowdsourcing

The term crowdsourcing firstly appeared in 2006 with the meaning of “taking a function once performed by employees and outsourcing it to an undefined (and generally large) network of people in the form of an open call” [6]. Over the years, this idea has gained much interest both in industry and academia and became an umbrella term that encompasses many different use cases, as distributed-human-intelligence tasks, knowledge discovery and management, etc. [3]. One of the most famous examples is Amazon Mechanical Turk, a crowdsourcing platforms where requesters can submit small jobs, such as image labeling, called Human Intelligence Tasks (HITs), which are performed by a remote community of workers. In recent years, researchers explored novel ideas to use crowdsourcing such as software development [11] and even writing fiction [7].

A persistent problem in crowdsourcing is how to engage crowd workers when performing HITs. Usually crowd workers are paid a small amount of money [10], which does not guarantee that workers perform to the best of their ability, even if a fair payment can be helpful [5], [15]. Another way of engaging workers is through gamification: in [1] a novel platform is designed, where users playing an online game were labeling images while competing against each other.

Providing engaged and motivated workers does not always guarantee that their answers will be correct. Therefore each HIT is usually assigned to multiple crowd workers, whose inputs are aggregated with statistical methods. The simplest way to aggregate multiple answers is by Majority Vote (MV), i.e., the answer that is the most popular among the workers is predicted to be correct. This, however, does not take into consideration differences in workers’ expertise.

Expectation Maximization (EM) algorithm has been used as early as in [4] to tackle this problem. EM iteratively estimates the probability that a worker correctly completes a set of HITs (this probability is called reliability in § IV). Some limitations of EM are that it needs to run offline on the whole dataset and it is costly due to its iterative nature. Both of these arguments prevent EM application in real-time scenarios.

A relatively new type of crowdsourcing is *Continuous crowdsourcing* [8] which engages workers over long periods and allows them to maintain the context of the task. Moreover the tasks change over time reacting to workers’ input. The responses from the workers are processed in so-called *input mediators* that aggregate them to produce the system output. Continuous crowdsourcing has been applied to a variety of use cases, for example, real-time continuous video/audio captioning to collect activity recognition labels from video stream [9], and a robot control system able to react to unforeseen circumstances [8].

Continuous crowdsourcing requires to dynamically estimate workers’ reliabilities. Such calculations need to be done in real-time, as workers’ input is immediately applied to the task at hand and different reliabilities affect the votes aggregation. In [13] a modified version of the EM algorithm is used in real-

time to estimate workers’ reliabilities. The authors, however, do not mention the time complexity of this approach, and using an iterative method in real-time scenarios might not be optimal. In this paper, we propose a method of calculating workers reliabilities in a real-time scenario without the need to perform high complexity calculations.

B. Video Games Controlled by Multiple Players

There is a growing trend of multiplayer video games, where multiple network users control a single instance of a video game streamed through an online platform. One of the first, and to this day the most popular event using this approach was Twitch Plays Pokemon, where thousands of players watched a video stream of an ongoing Pokemon Red game while issuing commands through the chat. The game creators defined two modes: *anarchy*, all inputs from all players are forwarded to the game; and *democracy*, all inputs are considered as votes, counted and then forwarded to the game, similarly to input mediators (§ II-A). The democracy mode resulted in a more conservative and slower behavior of the character, but allowed the players to beat the hardest challenges.

Twitch Plays Pokemon is one of the earliest examples of an Emergent Multiplayer Game (EMG) [14], a video game in which large crowds play collectively through a network voting system. EMGs utilize a set of relatively simple game mechanics and subsystems, that when controlled by large crowds result in creative and unique ways of playing. They encourage coordination and self-organization of the crowd, both in-game and outside of it. Explored mechanics usually use a voting system that gathers input from a large crowd and applies the results after a certain delay. The time reserved for collecting the input, however, makes the system not adequate for fast-paced, real-time scenarios. We address this challenge using our Crowdsourcing Controller.

III. SYSTEM OVERVIEW AND USE CASE

Our continuous crowdsourcing system is presented in Figure 1. The central part of the system is a simple video game (described in § III-A). Players issue inputs through the network as if they were the only player in control of the game character. The input is aggregated and processed by the *Crowdsourcing Controller*. The output of the system is a singular game command, whose type and issue time are calculated based on all the inputs sent over the network. The aggregated game command is applied with immediate effect, which influences the state of the game and is seen by all players on their screens. The source code of the implementation of our system is publicly available¹.

A. Use case: a Simple Video Game

Our Crowdsourcing Controller controls a character in a simple video game which we describe next. We choose this use case because: (1) crowd workers are generally familiar with simple controls as well as typical challenges presented in video games; (2) we wanted to explore the presented crowdsourcing

¹<https://github.com/KacperKenjiLesniak/crowded-dinosaur>

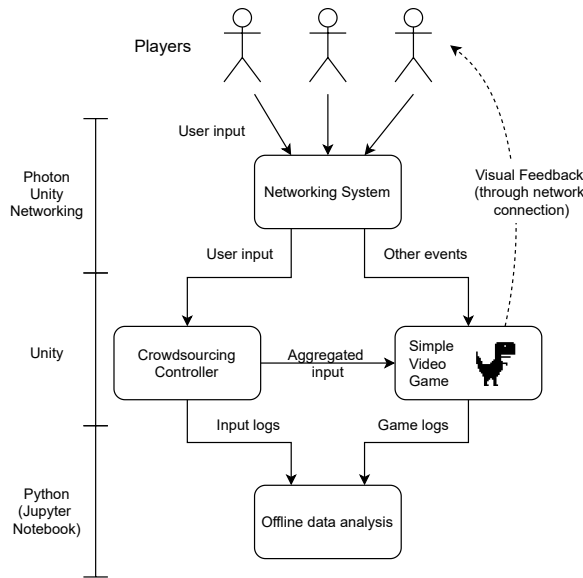


Fig. 1. General system overview with technologies used for implementation




Fig. 2. Crowd controlled dinosaur (black) and player dinosaurs (different colors) jumping over an obstacle

system as part of a novel gameplay mechanic; (3) gamification has proven to be a viable method to engage users in a variety of tasks [12] which can help with keeping workers motivated throughout the course of experiments.

The implemented game is a clone of a well-known mini-game played in Chrome (see Figure 2). Players are assigned their own dinosaur character of a different color. The crowd-controlled dinosaur which corresponds to the aggregated output is represented by a black dinosaur. Players' dinosaurs behave like "ghosts" and they serve purely as a visual clue to the players (Figure 2). Thanks to them the players can see that their input is being processed and is working as intended, even when the crowdsourced dinosaur is not reflecting their actions at the current moment. The players also see each other characters, which can enhance engagement and competition, and it provides further visual feedback. Players seeing the black dinosaur impacting an obstacle, even when they issued a correct input, can feel frustrated and believe that the system is not working. Seeing that others failed to issue the correct input can help alleviate this issue. Moreover, seeing all other characters can lead to players quickly realizing that they are

issuing the inputs out of sync.

We chose this game because it meets a specific set of requirements: the input comes in real-time and jumping or crouching at the right moment is crucial for progression; the input is discrete (short jump, long jump, crouch); there is a player score which can be used as an estimate of player skills, and the game is widely known and simple to play.

IV. CROWDSOURCING CONTROLLER

We describe the Crowdsourcing Controller next, which is responsible for analyzing the inputs coming from the players, aggregating them, and issuing a single command that best represents the combined inputs.

Let us consider a system where four players are able to issue two types of input - 1 and 2. The example in Figure 3 presents a certain moment in time, during which the system receives four inputs, each coming from a different player, three of them are of type 1 and one is of type 2.




Fig. 3. Inputs of 2 types issued by 4 players in a short period of time

This situation poses two challenges. The first challenge relates to the definition of the *input frame*, i.e., the time window where all inputs correspond to the players' reaction to the same game event, for example an obstacle appearing on screen. In the example in Figure 3, all 4 inputs are issued at a similar time, so it is reasonable to consider them in the same group or input frame. As for when the input frame should be analyzed further by the system, in the presented example it is reasonable to do it immediately after receiving the last input. Moreover, since no more inputs are issued after the last one, all inputs can be grouped and processed by the system for further analysis. However, it is not possible to make similar decisions in a real-time setting without knowing what comes next. Section IV-A describes how we tackle this problem with dynamic and static input frames.

The second challenge relates to the aggregation approach, i.e., how to combine multiple inputs in a single one. In the presented example, three players voted for input 1 and one player for input 2, so according to majority vote, the output of the system should be 1. However, majority vote treats all players equally and does not consider the expertise of each player, e.g., the vote of an expert or reliable player can weight more than the vote from a beginner and lead to a more accurate aggregated score. Section IV-B describes how we estimate such weights, called reliabilities.

The above mentioned challenges are handled by *input mediators*, illustrated in Figure 4. Input mediators are components of the crowdsourcing system that: (1) process the input stream coming in real-time from several players and group inputs into frames; and (2) produce the final aggregated output that is issued to the game. These two components together define our Crowdsourcing Controller.

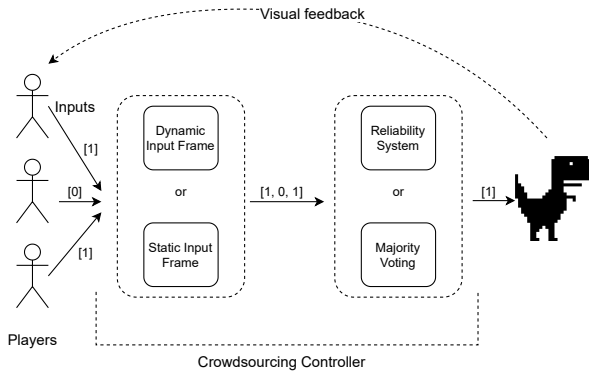


Fig. 4. Flow of players' inputs through different input mediators

A. Dynamic and static Input Frames

The first component of the Crowdsourcing Controller groups continuous input into input frames, which contain no more than a single input for each player (a missing input from a player is represented as 0). A frame represents a short moment in time in which the users issue the commands which can result in a specific behavior of the system. We use 2 types of input frames, static and dynamic, which we describe next.

The Static Input Frames [8] slice the input stream into frames of a constant length. In case of multiple inputs from the same player in a single frame, the system uses the most recent one. One limitation of this approach is how to define the frame length: too short frames will wrongly group inputs related to the same event in different frames, too long frames will negatively affect the system responsiveness.

To address this limitation, we propose the Dynamic Input Frame approach, where the frame length is dynamically defined based on players' inputs. The system forms a queue of incoming inputs and decides in real-time when to slice the input frame. The inputs in the queue have a specific Time To Live (TTL), after which they are discarded. When the number of inputs in the queue surpasses a threshold (T), the system begins the procedure of defining the frame. It waits for a time equal to half of the TTL, then gathers all the inputs in the queue, passes them to the aggregator component, and clears the input queue. It is crucial to correctly estimate TTL, as it creates a trade-off between the responsiveness of the system and its accuracy, as illustrated in Figure 5.

When the Dynamic Input Frame is used together with the Continuous Reliability System (§ IV-B), instead of using the number of inputs to compare to the threshold (T) we calculate a weighted sum using players' reliabilities, so that reliable players have more control over when the input frame is issued.

B. Inputs Aggregator

The second component of the Crowdsourcing Controller gets all the inputs within the input frame and returns a single aggregated input. We consider two approaches to aggregate inputs: (1) Majority Vote (MV); and (2) our Continuous Reliability System.




Fig. 5. Dynamic frames, threshold (T) of 2 players out of 4 marks the center of the frame: a) too short duration, part of input ignored; b) too long duration, unresponsive system; c) correct duration for the situation

MV simply selects the most popular input among players. Ties are arbitrarily broken with a predefined order of preferred inputs. Note that no or missing input also counts as a valid input, because sometimes the correct action is not issuing any commands at all.

The Continuous Reliability System uses weighted MV, where the votes are weighted according to the reliability score of each player, i.e., an estimate of the ability of each player to issue correct inputs. Next, we describe how to compute and update reliabilities in real-time.

Inspired by the idea of Collective Intelligence, the Continuous Reliability System rewards players that issue the same input as the majority and punishes those that vote for unpopular options. The system starts by aggregating inputs with MV, i.e., uniform reliabilities. Every time the system receives an input frame, the reliability scores are updated proportionally to how popular the players' votes are. For example, when 95% of the players vote for input A and 5% vote for input B, the reliabilities of the players that voted for option B will drop significantly, compared to a situation where a split between A and B is 60% to 40%. After a certain number of iterations, reliability scores might change the final output: the aggregated input is not the most popular one, but a less popular option can be chosen because of the high reliability of the players who issued that command. This system is completely decoupled from the game itself, it does not need to be aware of whether an input was *good* or *bad* based on the game state, which makes it easier to apply it to other scenarios in the future.

Next we formally describe the Continuous Reliability algorithm (Algorithm 1). Let $\mathcal{I} = \{0, 1, \dots, K-1\}$ be the list of all possible inputs issued by players and N_p be the number of players. A single input frame \mathcal{F} is an N_p -dimensional vector with values from \mathcal{I} , i.e., the votes of the players within the time frame. R_j is the reliability score of player j and R is the N_p -dimensional vector containing all current reliability scores.

The Continuous Reliability algorithm is recurrent. At the start, each player is assigned the same reliability score of 1, then at each iteration we compute ΔR , i.e., the vector of reliability changes that is added to update reliability scores.

Consider a fixed iteration with input frame \mathcal{F} . We define $\mathcal{V} = \{v_0, v_1, \dots, v_{K-1}\}$ as a K -dimensional vector counting the number of votes for each input in \mathcal{F} . We say that an input j is *viable* when it receives “enough” votes, i.e., $v_j \geq \gamma * v_{max}$, where γ is an *agreement threshold* and $v_{max} = \max(\mathcal{V})$, i.e., the maximum number of votes achieved by a single input. This handles cases when the votes are evenly split between more than one inputs: we do not want to decide that only one input is correct, but we will consider all viable votes when updating player reliabilities.

If all inputs in \mathcal{F} are viable we do not change reliabilities, because all players selected a potentially correct input. If there are non-viable inputs, we proceed to compute ΔR .

First we treat non-viable inputs in \mathcal{F} , intuitively they should affect ΔR in a proportional way: the less popular an input, the lower the player reliability who issued that. We define this as follows:

$$\Delta R_{j_{nv}} = -\frac{v_{max}}{v_{\mathcal{F}_j}} * \delta \quad (1)$$

where j is a player who issued a non-viable input; $v_{\mathcal{F}_j}$ is the number of votes received by the input issued by player j ; and $\delta \geq 0$ is the *reliability step coefficient*, which controls how much a viable or non-viable input affects the reliability score. The higher δ , the faster the system will adjust the reliabilities. The current reliability of player j , R_j , is updated by adding (subtracting) the quantity in Equation (1).

Note that the reliability scores might fall below 0, which means that unreliable players will in fact decrease the number of votes for the input they choose. This mechanism can be effective when dealing with players that issue wrong inputs on purpose, as they still unknowingly help the system in selecting the correct input. On the other side, a limitation of this mechanism arises when the same player keeps on issuing wrong inputs, e.g., inactivity, network lag, etc. If the player resumes the game and starts issuing correct inputs, his/her very low reliability will disrupt the voting, since weighted MV will heavily favor options not chosen by this player. Therefore, we introduce a *lower bound* $\omega \geq 0$, which clamps the change of the reliability if it brings the total player reliability below $-\omega$:

$$\Delta R_{j_{nv}} = \begin{cases} \frac{v_{max}}{v_{\mathcal{F}_j}} * \delta, & \text{if } R_j + \Delta R_{j_{nv}} \geq -\omega \\ \omega - R_j, & \text{otherwise} \end{cases}$$

where R_j is the current reliability score for player j . The opposite problem is not a concern for the upper bound. For players to have a high reliability score, they not only have to be consistent in issuing the input, but also other players have to vote similarly to them. This makes a situation where one of the players has a significantly larger reliability score than others very unlikely.

Then, we treat viable inputs in \mathcal{F} . Again they should affect ΔR in a proportional way: the more popular an input, the higher the player reliability who issued that. We do this as follows:

$$\Delta R_{j_v} = v_{\mathcal{F}_j} * \frac{\sum_{k \in nv} \Delta R_k}{\sum_{k \in v} v_{\mathcal{F}_k}} \quad (2)$$

Algorithm 1: Calculating the reliability updates

Input: Players input frame \mathcal{F} , current reliabilities
 $R = [R_0, \dots, R_{N_p-1}]$

Result: Vector of reliabilities updates
 $\Delta R = [\Delta R_0, \dots, \Delta R_{N_p-1}]$

$\Delta R = [0, 0, \dots, 0]_{N_p}$

$v_{\mathcal{F}_j} = \sum_k (\mathcal{F}_j == \mathcal{F}_k) //$ Number of votes for the input issued by player j

if $\forall j \in \text{viable}(v_{\mathcal{F}_j})$ **then**
| **return** ΔR

end

// First pass - non-viable votes

for *not viable*($v_{\mathcal{F}_j}$) $\in \mathcal{V}$ **do**

| $\Delta R_j = -\frac{v_{max}}{v_{\mathcal{F}_j}} * \delta$

| **if** $R_j + \Delta R_j \leq -\omega$ **then**

| | $\Delta R_j = \omega - R_j$

| **end**

end

$norm = \frac{\sum_{k \in nv} \Delta R_k}{\sum_{k \in v} v_{\mathcal{F}_k}} //$ Viable votes normalization factor

// Second pass - viable votes

for *viable*($v_{\mathcal{F}_j}$) $\in \mathcal{V}$ **do**

| $\Delta R_j = v_{\mathcal{F}_j} * norm$

end

return ΔR

where the first factor $v_{\mathcal{F}_j}$ is the number of votes received by the input issued by player j and the second factor is a normalization factor. The normalization factor is computed as the ratio of the sum ΔR for non-viable inputs and the sum of $v_{\mathcal{F}_k}$ for the viable inputs. The normalization factor is needed to guarantee that the sum of reliabilities remains constant and equal to the number of players, which is needed by weighted MV.

V. EXPERIMENTAL EVALUATION

Next we evaluate the performance of the Crowdsourcing Controller, i.e., input frames and input aggregators.

A. Experimental Setup

The Crowdsourcing Controller depends on a number of parameters. For static input frames we need to set only the frame length. For dynamic input frames we need to set the TTL and the frame threshold T , i.e., the number of inputs from different users needed to determine the centre of the dynamic time window (see Figure 5). We set TTL to 300ms, so on average the time between issuing an input and seeing a response on the screen is around 150ms. Threshold $T = 0.5$ was set so that half of the inputs marked the centre of the dynamic frame. For the continuous reliability system, we need to estimate the agreement threshold $\gamma \in [0, 1]$, i.e., the proportion of votes needed to define a viable input, and the reliability step coefficient $\delta \in [0, 1]$, which determines the speed in updating reliabilities. The parameters have been fine-tuned for offline and online experiments during development

and used consistently to compare the system under the same circumstances. The values were chosen so that the learning process of the system can stabilize in few minutes, as longer sessions could cause players’ fatigue. The values used were $\gamma = 0.8$ (offline), $\gamma = 0.6$ (online) and $\delta = 0.05$. The lower bound ω is set to 0.5.

For the input frame component, we consider the simple static input frame approach as our baseline opposed to our dynamic input frame approach. For the inputs aggregator component we consider MV as baseline opposed to our Continuous Reliability System. In the experiments we consider 4 different system configurations: (1) static frame with Continuous Reliability System; (2) Dynamic Input Frame with Continuous Reliability System; (3) Static frame with MV; (4) Dynamic Input Frame with MV.

For the evaluation, there is no available ground-truth, i.e., single correct sequence of inputs, that we can use to evaluate the Crowdsourcing Controller. This is because the game is dynamic and there is more than one way of beating obstacles, e.g., some obstacles can be beaten with both short and long jumps. Therefore we use the game score as evaluation measure. The game score represents how long it takes for the players to lose the game, by issuing the wrong type of input or providing it at the incorrect time. The higher the game score the better.

Additionally we consider EM [4] as our reference or “ground-truth”. Due to its iterative definition, EM can not be run online, therefore it is run offline when all data are gathered. Then we compute the agreement between one of the 4 configurations and EM, i.e., the accuracy with the inputs issued by a configuration and the inputs predicted by EM as ground-truth. The higher the agreement the better.

B. Offline Experiment with Synthetic Players

In this experimental scenario we evaluate the Crowdsourcing Controller in an artificial setting with synthetic players. Players are replaced with AIs designed to represent humans of different skills. The *Perfect AI* always issues a correct input at each time-stamp during the game. All other AIs are defined by adding noise to the perfect AI behaviour. This is done with 3 parameters representing: noise, noise shift and chance of error. We generated a total of 14 different AIs: 1 Perfect, 3 Good, 4 Bad, 1 Shifted (-), 2 Shifted (+), 1 Majorly Shifted (-), 1 Confused, and 1 Random. Each AI was individually tested 3 times and their mean scores are in Table I.

The comparison of scores achieved by different system configurations across 6 consecutive trials is presented in Figure 6 and in Table II. We stopped after 6 trials since we reached convergence of the reliability scores, i.e., after trial 6 there was little change in reliability scores.

From the results we can see that the Dynamic Input Frame configurations outperform the static ones, both in terms of game score and agreement with EM. Moreover, the configurations with the Continuous Reliability system outperformed MV, both in terms of scores and similarity with respect to

TABLE I
AIs INDIVIDUAL GAME SCORES

Name	Average Score	Max Score
Perfect	1182	1353
Good	250	488
Bad	63	66
Confused	23	28
Random	48	64
Shifted (-)	164	398
Shifted (+)	27	30
Majorly Shifted (-)	98	175




Fig. 6. 14 AIs scores with different system configurations

EM. For the Dynamic Input Frame with Continuous Reliability System (the best configuration), the agreement with EM reached a value of 99.37% compared to 94.40% if the system is replaced with majority voting (5.26% relative improvement). In Figure 6, the Dynamic Input Frame with Continuous Reliability System outperforms all other configurations already after 1 trial. This means that the system quickly and correctly learns which players are more reliable.

Figure 7 shows the process of the reliability training with the Dynamic Input Frame configuration. The reliability scores (y -axes) are plotted for each of the analyzed input frames (x -axis). The process is chaotic at the beginning but stabilizes after circa 100 frames. After 100 frames, the estimated reliabilities correlate with the AIs skills: the top reliability score (red line) belongs to the Perfect AI. The inputs issued by the perfect AI will be weighted as 7 players. The following 3 reliabilities scores belong to the 3 Good AIs. At the end of the experiment, the Perfect AI with the help of just one of the Good AIs can outvote the rest 12 AIs combined. Indeed, the top 4 AIs alone have an impact of over 18 players. This

TABLE II
14 AIs EXPERIMENT RESULTS

Name	Aggrement w/ EM predictions	Mean score	Max score
Dynamic Reliability	0.9937	852	1211
Dynamic Majority	0.9915	474	981
Static Reliability	0.9553	309	610
Static Majority	0.9714	94	277




Fig. 7. 14 AIs reliability training with the Dynamic Input Frame

TABLE III
USER STUDIES RESULT SUMMARY

Name	Aggrement w/ EM predictions	Mean score	High score
Dynamic Reliability	0.9619	255	541
Dynamic Majority	1.00	194	516
Static Reliability	0.9108	76	434

exceeds the total number of players (14) because the rest of the reliabilities are less than 0. We can also observe the effect of the lower bound ω , which prevents bad players from disrupting the weighted voting system.

C. Online Experiment with Real Players

In this experimental scenario, we evaluate the proposed Crowdsourcing Controller with real users. We tested only 3 configurations out of 4: Dynamic Input Frame with MV, Dynamic Input Frame with Continuous Reliability System, and static frame with Continuous Reliability System. We did not test the static frame with MV because: (1) the experiment with 4 configurations was too long and players' fatigue could affect the quality of the collected data; (2) the offline experiments show that this is the worst performing configuration (see § V-B)

We enrolled a total of 9 players, all of them were university students with a scientific background. Before starting the experiment, participants were trained by playing individually for 3 minutes. Afterward they had 3 trials to achieve their highest scores. Then we Started the multiplayer phase. Each configuration was played for 3 minutes for training, followed by 5 trials to beat the highest score. The whole process was repeated 2 times to collect more data. Exact user instructions can be found in our code repository. During the experiments the reliability scores were not shown to the players. This was done to avoid the possibility of the players feeling discouraged due to low impact of their inputs. There was a short break between the sets of configurations to give the players a chance to rest.

The game scores achieved by the players during the first set of trials (both training and game) are shown in Figure 8. The

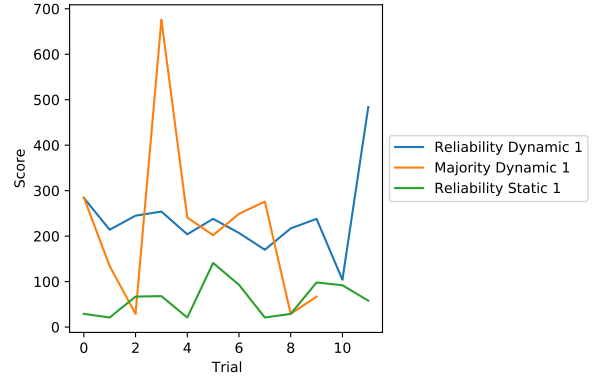


Fig. 8. User scores during first set of trials

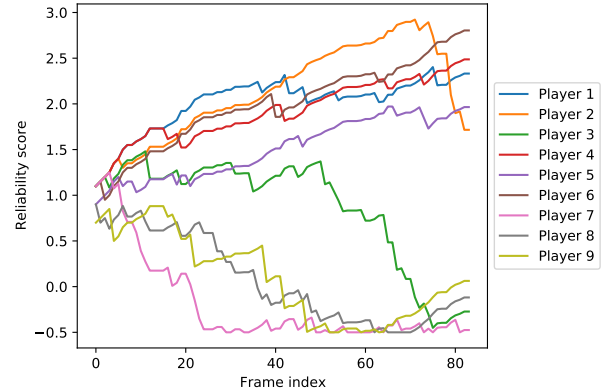


Fig. 9. Player reliability training in one of the sessions with the Dynamic Input Frame

results for the second trial are comparable. Table III reports the results aggregated over all trials.

Online experiments corroborate the observations from offline experiments with respect to input frames. The Dynamic Input frame configurations outperform the static input frame configuration (green line in Figure 8). This is true both in terms of game score and agreement with EM. Second, in terms of game score, the Dynamic Input Frame with Continuous Reliability outperforms the other 2 configurations.

Online experiments differ from offline experiment with respect to the agreement with EM. We can observe that using a real-time MV system with dynamic frames yields the same results as using an offline-trained EM. This suggests that the majority of the players are similarly skilled, thus using a weighted MV instead of the uniform MV does not make a significant difference. We further analyzed the agreement among the players which showed that a particular majority (specifically 5 players out of 9) agreed with each other on issuing the same input over 80% of the times. This can be seen also in Figure 9, where the reliability scores of the top 5 players are clearly separated from the rest of the group.

Even if all the players are similarly skilled, the better performance of the Continuous Reliability system in terms of

game scores might come from the fact that fewer players with stronger influence more quickly surpass the needed threshold of votes required by the Dynamic Input Frame, which helps with beating obstacles that require fast decisions.

At the end of the user studies, participants filled a short survey. Among different observations, over half of the participants mentioned, in different words, that they focused more during multiplayer than individual sessions, which is aligned with previous results about gamification approaches to encourage crowd workers (see Section II-A). Some participants mentioned the social aspect of playing together, saying that winning and losing together was more *fun*. This suggests that the system could be used as a gameplay mechanic in EMGs. Some players, however, reported feeling frustrated when they lost despite seemingly issuing correct inputs, sometimes blaming it on the network delay or software errors. This implies that players could benefit from even more visual feedback regarding the decision making of the Crowdsourcing Controller.

D. Discussion

Next, we summarize the main conclusions and observations from offline and online experiments. The Dynamic Input Frame proves to be more efficient than a static approach proposed earlier in the literature. It is able to achieve much higher scores thanks to more precise predictions of the correct time that the input should be issued at.

The Continuous Reliability System increases the game score to a large extent for offline experiment, but to a smaller extent for online experiments. The reason for this is the observed lack of differences in the skills of real players. Although a small number of them did in fact perform worse than others, the majority usually agreed on the input type and time. In this situation, it is expected that MV gives similar results as the Continuous Reliability System. Moreover, analyzing players' inputs proves that participants usually agreed on the type of the input, with the differences mostly showing in the precise time of issuing it. The reason for this comes from the nature of the game that was used to test the system. The real challenge was in fact issuing the input at the right time, since choosing the right input was relatively simple as different obstacles were easily distinguishable.

Both user surveys and observations during the study showed a significant social aspect of playing the game together, where synchronization seemed like one of the goals that the players wanted to achieve, knowing it would raise their chances of getting high scores.

VI. CONCLUSIONS AND FUTURE WORK

In this paper we present a prototype of an online video game, where multiple players control simultaneously the same character. We propose the Crowdsourcing Controller, i.e., a system that: (1) groups real-time inputs from multiple players into time frames; (2) aggregates the inputs within a frame and returns a single output command. We propose the Dynamic Input Frame method, which adjusts the frame length

dynamically based on the inputs received. To aggregate the inputs, we propose the Continuous Reliability System, which computes a weighted MV by estimating players' reliabilities. Offline and online experiments show that the Dynamic Input Frame is better than state-of-the-art static frames and that the Continuous Reliability System leads to higher game scores.

As future work we plan to test the proposed crowdsourcing system with a larger and more diverse set of players. Moreover, we will explore how the proposed systems cope in different settings, for example scenarios featuring more agency and creativity on the players' side. Finally, we want to inspect the possibilities of automatically optimizing some of the system parameters that were arbitrarily chosen, for example the TTL of the Dynamic Input Frame.

REFERENCES

- [1] Ahn, L. & Dabbish, L. Labeling Images with a Computer Game. *Proceedings Of The SIGCHI Conference On Human Factors In Computing Systems*. pp. 319-326 (2004), <https://doi.org/10.1145/985692.985733>
- [2] Anderson, S. Watching people is not a game: Interactive online corporality, Twitch.tv and videogame streams. *Game Studies*. **17** (2017,7)
- [3] Brabham, D. Crowdsourcing. (The MIT Press,2013)
- [4] Dawid, P., Skene, A., Dawid, A. & Skene, A. Maximum likelihood estimation of observer error-rates using the EM algorithm. *Applied Statistics*. pp. 20-28 (1979)
- [5] D'Eon, G., Goh, J., Larson, K. & Law, E. Paying Crowd Workers for Collaborative Work. *Proc. ACM Hum.-Comput. Interact.* **3** (2019,11), <https://doi.org/10.1145/3359227>
- [6] Howe, J. The Rise of Crowdsourcing. *Wired*. (2006,6), <https://www.wired.com/2006/06/crowds/>
- [7] Kim, J., Serman, S., Cohen, A. & Bernstein, M. Mechanical Novel: Crowdsourcing Complex Work through Reflection and Revision. (2016,11)
- [8] Lasecki, W., Murray, K., White, S., Miller, R. & Bigham, J. Real-Time Crowd Control of Existing Interfaces. *Proceedings Of The 24th Annual ACM Symposium On User Interface Software And Technology*. pp. 23-32 (2011), <https://doi.org/10.1145/2047196.2047200>
- [9] Lasecki, W. & Bigham, J. Interactive Crowds: Real-Time Crowdsourcing and Crowd Agents. *Handbook Of Human Computation*. pp. 509-521 (2013,11)
- [10] Silberman, M., Tomlinson, B., LaPlante, R., Ross, J., Irani, L. & Zaldivar, A. Responsible Research with Crowds: Pay Crowdworkers at Least Minimum Wage. *Commun. ACM*. **61**, 39-41 (2018,2), <https://doi.org/10.1145/3180492>
- [11] Stol, K. & Fitzgerald, B. Two's Company, Three's a Crowd: A Case Study of Crowdsourcing Software Development. (2014,5)
- [12] Tomaselli, F., Sanchez, O. & Brown, S. How to Engage Users through Gamification: The Prevalent Effects of Playing and Mastering over Competing. (2015,12)
- [13] Welinder, P. & Perona, P. Online crowdsourcing: Rating annotators and obtaining cost-effective labels. *2010 IEEE Computer Society Conference On Computer Vision And Pattern Recognition - Workshops*. pp. 25-32 (2010)
- [14] Wodarczyk, S. & Von Mammen, S. Emergent Multiplayer Games. *2020 IEEE Conference On Games (CoG)*. pp. 33-40 (2020)
- [15] Ye, T., You, S. & Robert Jr, L. When does more money work? Examining the role of perceived fairness in pay on the performance quality of crowdworkers. *Proceedings Of The International AAAI Conference On Web And Social Media*. **11**, 327-336 (2017)