
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Roohi, Shaghayegh; Relas, Asko; Takatalo, Jari; Heiskanen, Henri; Hämäläinen, Perttu
Predicting Game Difficulty and Churn Without Players

Published in:
CHI PLAY 2020 - Proceedings of the Annual Symposium on Computer-Human Interaction in Play

DOI:
[10.1145/3410404.3414235](https://doi.org/10.1145/3410404.3414235)

Published: 02/11/2020

Document Version
Peer-reviewed accepted author manuscript, also known as Final accepted manuscript or Post-print

Please cite the original version:
Roohi, S., Relas, A., Takatalo, J., Heiskanen, H., & Hämäläinen, P. (2020). Predicting Game Difficulty and Churn Without Players. In *CHI PLAY 2020 - Proceedings of the Annual Symposium on Computer-Human Interaction in Play* (pp. 585-593). ACM. <https://doi.org/10.1145/3410404.3414235>

Predicting Game Difficulty and Churn Without Players

Shaghayegh Roohi
shaghayegh.roohi@aalto.fi
Aalto University
Espoo, Finland

Asko Relas
asko.relas@rovio.com
Rovio Entertainment
Espoo, Finland

Jari Takatalo
jari.takatalo@rovio.com
Rovio Entertainment
Espoo, Finland

Henri Heiskanen
henri.heiskanen@rovio.com
Rovio Entertainment
Espoo, Finland

Perttu Hämäläinen
perttu.hamalainen@aalto.fi
Aalto University
Espoo, Finland

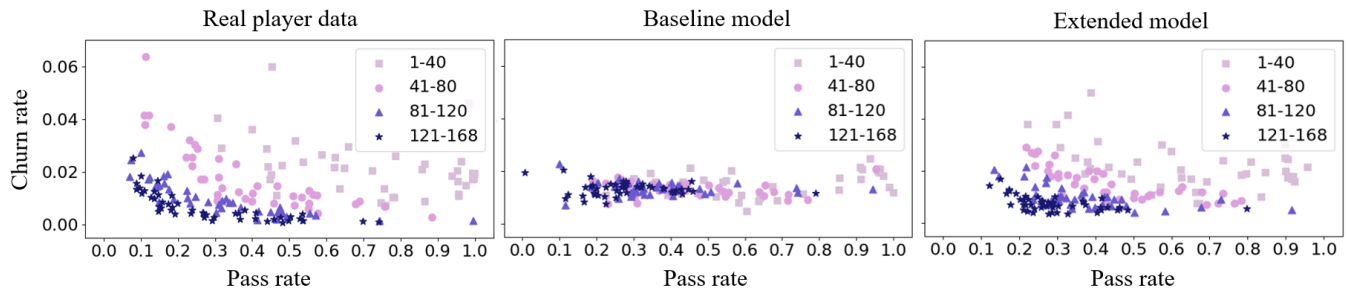


Figure 1: Scatter plots depicting the relation of pass rate (a measure of level difficulty) and churn rate over 168 game levels of Angry Birds Dream Blast, in both real player data and our simulations. Here, churn is defined as not playing for 7 days. The colors denote level numbers. The baseline simulation model predicts pass rate and churn directly from AI gameplay. Our proposed extended model augments this with a simulation of how the player population evolves over the levels.

ABSTRACT

We propose a novel simulation model that is able to predict the per-level churn and pass rates of Angry Birds Dream Blast, a popular mobile free-to-play game. Our primary contribution is to combine AI gameplay using Deep Reinforcement Learning (DRL) with a simulation of how the player population evolves over the levels. The AI players predict level difficulty, which is used to drive a player population model with simulated skill, persistence, and boredom. This allows us to model, e.g., how less persistent and skilled players are more sensitive to high difficulty, and how such players churn early, which makes the player population and the relation between difficulty and churn evolve level by level. Our work demonstrates that player behavior predictions produced by DRL gameplay can be significantly improved by even a very simple population-level simulation of individual player differences, without requiring costly retraining of agents or collecting new DRL gameplay data for each simulated player.

CCS CONCEPTS

• **Human-centered computing** → User models; • **Computing methodologies** → Modeling and simulation.

KEYWORDS

Player Modeling; Game AI; Churn Prediction

ACM Reference Format:

Shaghayegh Roohi, Asko Relas, Jari Takatalo, Henri Heiskanen, and Perttu Hämäläinen. 2020. Predicting Game Difficulty and Churn Without Players. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play (CHI PLAY '20)*, November 2–4, 2020, Virtual Event, Canada. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3410404.3414235>

1 INTRODUCTION

One of the primary difficulties of game design and development is that player behavior is hard to predict. This leads to an iterative design process of prototyping and testing, which is slow and expensive. Ideally, research should produce models and tools that allow evaluating the effect of design decisions early on, before committing resources to real-life game testing. This is one of the foundational motivations of player and user modeling [29, 31, 49].

Better models and tools are in particular needed for predicting and optimizing business critical behavior such as churn, i.e., a player quitting the game and not coming back to it. Churn matters as many modern games accumulate their revenue gradually from in-game advertisements and purchases, instead of single up-front fee. To prevent churn, free-to-play game companies engage in extensive

data-driven A/B-testing and optimization of game levels. For example, one would like to identify and modify levels with high churn rate. However, instead of deploying different game level versions to real players and seeing what happens, it would be desirable to have predictive models of churn that allow designers to conduct initial testing and prototyping *in silico*.

An emerging possibility is to use simulated game-playing agents instead of human playtesters [3, 31], at a fraction of the time and cost. However, present models are limited, e.g., by the need to train models with large amounts of real player data in a way that does not necessarily generalize to new game content [11]. Models and algorithms that can play games without a large dataset of ground truth behaviors are often not validated to give realistic predictions of player behavior and experience [36].

This paper makes the following contributions to player modeling and simulated game testing:

- (1) We present new data indicating a relation between game level difficulty (measured as pass rate) and churn rate, shown in Figure 1. The data is from 168 levels and 95266 players of Angry Birds Dream Blast, a successful free-to-play mobile game from Rovio Entertainment.
- (2) We propose a novel simulation model that predicts the observed relation of pass and churn rates. Our key innovation is to combine Deep Reinforcement Learning (DRL) game playing agents with a simulation of how a player population with simple computational models of skill, persistence, and boredom evolves over the levels of a game as some players churn.

Previous work has investigated DRL gameplay [16] and player population simulation [35], but not combined the two. Previous work has also predicted pass rate [11] and churn [6, 32, 48], and found a relation between game success rate and engagement [24, 25], but has not modeled how the two interrelate in a dynamic manner over a game's levels.

2 RELATED WORK

2.1 Churn prediction

Churn means inactivity duration, and its precise definition varies from a few days to completely quitting a game. Churn prediction empowers game companies to engage players who are likely to churn, and hence increase the success probability of the game. Churn prediction has been approached through survival ensembles and Cox regression to estimate the likelihood of not being churned and survive after a specific amount of time [5, 32]. Bertens et al. [5] inspected both survival ensembles and Cox regression for using in-game features such as player logins, play time, purchases, and level ups in churn prediction.

Other methods address churn prediction as a classification problem. Bonometti et al. [6] employ deep neural networks to jointly estimate survival time and churn probability by modeling early interactions between players and the game, using the metrics like play time, time difference between play sessions, and in-game activity type and diversity. Yang et al. [48] study the regularity of the time that long-term players spend in the game to perform the binary classification into churned and not-churned.

In the context of churn classification, some works use aggregated data [48], and others exploit temporal data [19, 22] instead. Bonometti et al. [6] compared aggregated and temporal data, and find that models with temporal data outperform the other kind. Kristensen et al. [20] propose using stacked LSTM networks with a combination of aggregated and time-series data as their inputs.

In contrast to the work above, we do not predict the probability that a particular player churns. Instead, we predict the average churn ratio per game level, with the goal of helping game designers identify problematic levels before they are deployed to players. We employ game-playing AI agents in producing the input features of the predictor.

2.2 Game-playing Agents

Several studies have investigated game-playing agents in video games. Some of these studies have focused on Monte Carlo tree search (MCTS) [4, 13, 34], and others have exploited advances in deep learning methods [11, 17, 21, 26, 46]. For in-depth overviews of recent research on the topics, we refer the reader to Shao et al. [40] and Justesen et al. [16].

Game-playing agents are beneficial in playtesting by reducing costs and the need for human playtesters [3]. AI agents have been found useful in finding bugs [8], game parameter tuning [14], and predicting level difficulty and average player pass rate [11, 34].

Agents that are used in playtesting need to be reasonably human-like. Zhao et al. [50] defines human-likeness as a balance between skill and playing style. Ariyurek et al. [3] compare MCTS and human-like AI agents in finding bugs. One can train human-like game-playing agents simply through deep neural networks and imitation learning [11], although this requires large amounts of human gameplay data and does not necessarily generalize to new content. For instance, if there are no human reference actions to imitate for a new type of a game puzzle, an imitation learning agent's behavior is undefined.

An alternative to imitation learning is provided by the computational rationality paradigm, which views human behavior as emerging from optimization of rewards received through taking actions, albeit with limited computational capabilities [10]. The rewards can be either extrinsic ones such as a game score, or intrinsic ones defined by computational motivation and emotion models [36]. A common way to implement such reward optimization is through MCTS, which requires no lengthy training, but incurs a high run-time computing cost. For instance, Poromaa et al. [34] uses a MCTS agent's performance to predict the average human pass rate. An alternative is Deep Reinforcement Learning (DRL), which requires a lengthy and potentially unstable training process, but can yield computationally lightweight neural network agents. DRL is also better suited for continuous state and action spaces, which are needed for embodied physically simulated players. DRL has recently been shown to enable user simulation including the biomechanics of the human body, predicting the perceived shoulder fatigue of pointing movements [7].

In addition to imitation learning and computational rationality approaches, some works implement heuristic agents based on the common strategies that players take during playing the game [8, 13]. Isaksen et al. [14] optimize game parameters to reach a specific

level of difficulty based on a heuristic AI player’s score. Silva et al. [8] suggest how analysis of AI gameplay can be helpful in finding defects in the rule set of a board game. They explored various aspects of the game by matchups between game-specific agents, each with different playing style. More recently, Holmgard et al. [13] evaluate the effect of enhancing MCTS agents with player type heuristics on the playability of multiple dungeon maps.

In this paper, we focus on DRL agents that require no hand-coded heuristics or human player data for learning to play a game. Human data is only required for fitting a simple “sim2real” model — e.g., linear regression — that maps simulation results to human-like churn and pass rates.

2.3 Understanding Game Difficulty

Our work is motivated by the relation of level pass rate and churn rate observed in our data. Pass rate is a measure of level difficulty, and the effect of difficulty on player experience is a topic with a considerable body of research. One of the most prevalent models in game design literature is the flow channel [38, 43], according to which the ideal game is not too easy and not too hard, to keep the player in the sweet spot between boredom and anxiety. A similar view is provided by intrinsic motivation theories: Self-determination theory posits the basic psychological need to feel competent [37], which a too difficult game can obviously destroy. On the other hand, research has also identified a basic need for novelty of stimuli and experiences, central to experiencing curious interest [18, 25, 41]. Competence and curiosity interact, as one way to keep a game easy is to avoid introducing new types of challenges the player has to learn, which violates the need for novelty. In practice, successful games seem to optimize for both competence and curiosity, using a carefully tuned difficulty progression that allows experiencing competence, and periodically resets back to a low difficulty when introducing new mechanics and challenges to maintain novelty [23]. Recent work on mathematically optimal learning, e.g., in artificial neurons, also posits a flow-channel style difficulty model [47], identifying 85% success rate as optimal.

Naturally, preferences for difficulty vary between players and games. For example, recent research has provided explanations on why frustratingly difficult games can sometimes be enjoyable [33]. However, our data is more in line with Lomas et al. [24], who found lower difficulty to yield the highest engagement in their educational game. In a follow-up study, Lomas et al. [25] found moderate difficulty as most motivating when players could select their opponent’s skill ranking. However, the easiest games were most motivating when difficulty was determined randomly.

3 DATA: ANGRY BIRDS DREAM BLAST

This paper focuses on Angry Birds Dream Blast [9], a successful free-to-play mobile game. Angry Birds Dream Blast is a non-deterministic, physics-based match-3 game with a limited number of moves per level. The player pops adjacent bubbles with similar colors. Figure 2 shows a screenshot of the game. Beyond simple bubble popping, different types of boosters can be produced by combining bubbles. These boosters are used to pop a large quantity of bubbles in rows and columns. Boosters are required to remove some objects in the game. Additionally, there are locks in the game

that can be unlocked by collecting bubbles or other objects. Players can also acquire power-ups and extra moves through in-app purchases. However, this is a feature we did not enable for the AI agents.



Figure 2: Screenshot of Angry Birds Dream Blast.

For our study, we utilized a dataset of per-level pass and churn rates from a total of 95266 players. Per-level churn rate, in range 0...1, was defined as the portion of players who stopped playing for at least 7 days after trying the level at least once. After the 7 day period, some players may return to the game, but churn measurement time needs to be limited in practice, e.g., for A/B testing of game content updates. Pass rate, also in range 0...1, was computed as the mean of 1 divided by the number of attempts players required to pass each level.

Figure 1 indicates a relation between real player pass and churn rates. Although the overall correlation is small (Spearman’s $r = -0.144$), the correlation is large when computing churn rate of a level as the portion of players who churn without completing the level (Spearman’s $r = -0.586$). In other words, players who churn without completing a level are more likely to do so in the more difficult levels. Beyond a simple correlation, Figure 1 also suggests that a low pass rate causes less churn in the later levels. It is this dynamic we set out to explain in this paper.

It should be noted that Angry Birds Dream Blast is a live game that is constantly evolving. Our data was collected from a version that does not have all the mechanics of the latest game version.

4 METHOD

We implement and compare two churn and pass rate prediction models:

- (1) A baseline regression model that directly predicts churn and pass rate from AI gameplay, motivated by earlier work on predicting game difficulty by AI gameplay [11] and the relationship between pass rate and churn observed in our data in Figure 1, echoed by previous research on game difficulty and engagement [24, 25].
- (2) The extended model in Figure 3 that combines both AI gameplay and player population evolution over game levels.

The per-level pass and churn simulation of the extended model takes in both a level difficulty estimate based on the AI gameplay, and a player population. The simulation outputs both the churn and pass rate predictions of a level, and the new player population for the next level, with churned players removed from the population. In determining the passed and churned players, we employ simple computational models of the following psychological and learning phenomena:

- A player passes a level if the level difficulty is lower than the player’s skill.
- A player churns if taking more attempts to pass a level than the player’s persistence. This allows us to model how less persistent players churn earlier, leading to later levels being played by more persistent players who are less sensitive to churn because of high difficulty.
- Even if a player passes a level, they may churn with some random probability proportional to the player’s tendency to get bored.
- Passing a level gets gradually easier over multiple attempts, as the player learns from their mistakes.

Below, we first describe our AI game playing agent, then explain how the AI gameplay data is used for the baseline model. Finally, we detail the extended model of Figure 3, including the pass and churn simulation (Algorithm 1).

4.1 AI Gameplay

Our Deep Reinforcement Learning AI agents are implemented using the Unity ML-agents framework [15] and Proximal Policy Optimization (PPO) [39] algorithm.

In Reinforcement Learning, an agent observes system state and takes an action, which leads it to receive a new observation and a reward. PPO optimizes the policy — a distribution of actions conditional on state observations, implemented as a neural network — to maximize the agent’s expected cumulative future rewards.

To make game playing emerge with PPO, one needs to frame the game as a RL problem by defining the state observations, actions, and rewards. Furthermore, one needs to define the neural network architecture and a number of algorithm parameters. Our choices are detailed below.

4.1.1 State observations. A combination of visual and numerical vector observation were used. The visual observation is a $84 \times 84 \times 3$ RGB screenshot of the game, and the vector observation contains the number of moves left, the types and numbers of the remaining level goals (e.g., collecting a specific number of bubbles), the numbers and types of the objects needed to unlock the visible locks, and game camera position.

4.1.2 Action space. We use discrete action space, discretizing the game screen into 32×32 possible points that the agent can click/tap. A continuous action space was also tested, but did not produce as good results.

4.1.3 Reward function. The reward function combines both extrinsic rewards and an intrinsic reward generated by the self-supervised curiosity model of Pathak et al. [30] implemented in Unity ML-agents.

The extrinsic reward is computed as a sum of reward components: win bonus or lose penalty, cleared goals ratio, progress percentage in opening locks, a small constant negative reward that penalizes the agent for using more moves than needed, and a click reward. The click reward is calculated as $r_{distance} = c_0 \exp\left(-\frac{d}{c_1}\right)$, where d is the distance to the closest available match, and c_0 and c_1 are tuning parameters. During the initial learning exploration where the agent has no idea where to click, the click reward provides extra guidance towards only clicking on valid targets.

4.1.4 Training. We train the agent separately on each level for 5 million PPO iterations, which takes on average 60 hours per game level. Multiple training runs were conducted in parallel using cloud computing (Amazon Web Services [2]). Each run was performed on an AWS instance with a 16-core 3.4 GHz Intel Xeon CPU, and used approximately 10 Gb RAM.

We provide the agent with 4x more moves per level than available for human players, so that the positive rewards for completing a level are less sparse. Too sparse rewards are a common reason for RL agents not learning. The extra moves also help if the level difficulty is too high and the number of moves is not sufficient even for human players.

4.1.5 Neural network. Visual observations are encoded using a convolutional neural network, with the result concatenated with the other observations as shown in Figure 4. This is one of the default settings provided in Unity ML-agents, based on the architecture originally proposed by Mnih et al. [27].

4.1.6 PPO and Unity ML parameters. The iteration experience budget of PPO agent is 10240 game moves. The agent collects experiences in episodes with a large time horizon of 1024, i.e., each episode is played until the level is completed or the agent runs out of moves. We use optimization batch size 1024, decaying learning rate initialized at 0.0003, PPO clipping parameter of 0.2, PPO entropy coefficient 0.005, discount factor $\gamma = 0.99$, Generalized Advantage Estimation $\lambda = 0.95$, and curiosity reward coefficient 0.02.

The parameters are explained in detail in the original PPO paper [39] and Unity ML-agents [15]. The maximum number of moves per level multiplier and the curiosity reward coefficient were tuned manually. Other parameters use the Unity ML-agents default values.

4.2 Baseline model

The baseline model predicts both churn and pass rate using a simple least squares linear regression of the form $\mathbf{x}^T \mathbf{w} + b$, where \mathbf{x} is a feature vector for the game level of interest, \mathbf{w} is a vector of regression weights, and b is a bias term. The feature vector comprises the mean, standard deviation, min, max, and the 5th, 10th, 25th, 50th,

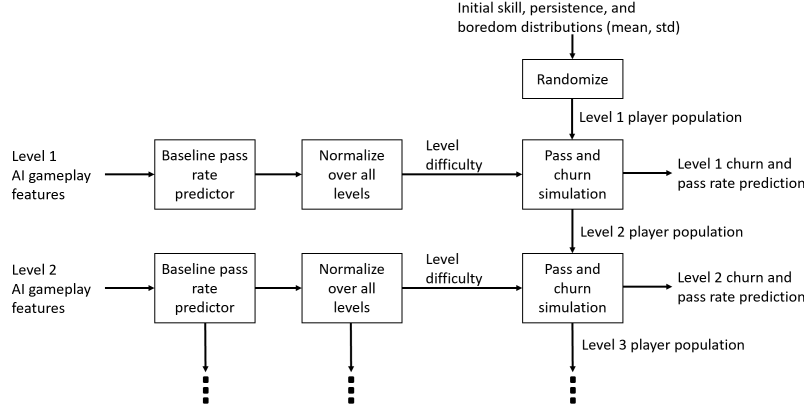


Figure 3: Overview of our extended simulation model, showing how AI gameplay data of each level results in a level difficulty estimate, which is fed to pass and churn simulation, together with player population parameters. The simulation outputs both the churn and pass rate predictions of a level, and the player population for the next level.

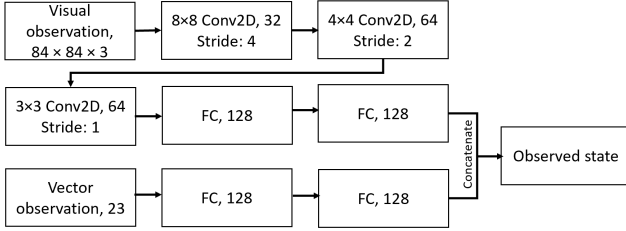


Figure 4: Neural network architecture of the observation encoder, combining both convolutional and fully connected (FC) layers.

and 75th percentiles of the cleared goals percentage while using the moves available for human players, and the mean, standard deviation, and the 5th, 10th, and 20th percentiles of the amount of moves left when passing the level, and the mean and standard deviation of the AI agent’s pass rate while using the same number of moves as human players. In total, this yields 16 features.

Figure 5 shows the relation between AI and human pass rate over levels. AI pass rate alone is not a great predictor of human pass rate, especially in the later game levels, where AI is only rarely able to pass the levels with the number of moves allowed for human players. As explained in section 4.1.4, the AI can use 4x more moves than human players. However, other AI gameplay statistics like cleared goals percentage and moves left ratio help in the prediction.

4.3 Extended Model

A basic problem of the baseline model is that it doesn’t take into account how the player population varies over the levels. Because the gameplay RL formulation does not include individual differences of players, and because the AI is trained on all levels, the difficulty predicted by the AI is an average measure that cannot map well to the difficulty experienced by players with different skills and personalities. Individual differences could be implemented as reward

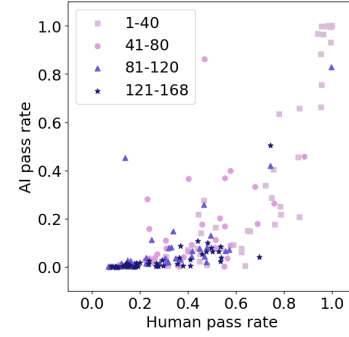


Figure 5: Relation between AI and human pass rate. Scatter plot colors correspond to game level indices.

function terms determined by intrinsic motivation and emotion models [36], but this would incur a considerable computational cost as each different agent would require its own lengthy training process. This is why our extended model simulates individual differences using a second, computationally much more simple population-level simulation model.

As shown in Figure 3, our extended model uses the baseline pass rate predictor as a building block, in conjunction with a pass and churn simulation block that also takes in the current player population, and outputs the population remaining after some players have churned. This way, the distribution of player skills and other attributes is allowed to evolve throughout the game’s levels. The DRL gameplay data needs to only be collected once, after which the population simulation can be run with different parameters with a very low computing cost. One simulation through all 168 levels takes less than a second on a standard desktop computer.

4.3.1 Pass and Churn Simulation Algorithm. The pass and churn simulation is detailed in Algorithm 1. We simulate a population of 2000 players, each player described by a tuple of real-valued attributes for skill, persistence, and boredom tendency. The initial

player population’s attributes are sampled from a normal distribution.

For each level, we draw two random numbers s, t from normal distributions with means determined by player skill and persistence attributes, and standard deviations defined by the parameters α, β (Algorithm 1 lines 5-6). Each simulated player attempts a level until passed or churned (line 9). If the random draw s exceeds the level difficulty, the player passes (line 11). Otherwise, the player churns if the number of attempts exceeds the random draw t (line 22). Additionally, a player also churns if passing the level but becoming bored, which is determined by the third random draw b based on the boredom attribute (lines 14-15). For each failed attempt, the players are also learning from their mistakes, simulated as incrementing s (line 21). At the end, we return the population to its original size by replicating randomly selected players (line 30). Because of the random selection, this replication does not affect the distribution of individual differences, and it prevents population depletion that would lead to inaccuracies in simulating later levels.

Note that we only simulate per-level learning, instead of learning over the levels, as our level difficulty estimates and ground truth data already have learning “built-in”. The DRL agent is trained from scratch for each level, and the human ground truth pass rates are measured from players who have played through the level progression, learning on the way.

Figure 6 shows how skill, persistence, and boredom attribute distributions change over levels in the simulation. One can observe an increase in average persistence, and a decrease in the tendency to get bored. There is also an increase in skill, although to a lesser degree.

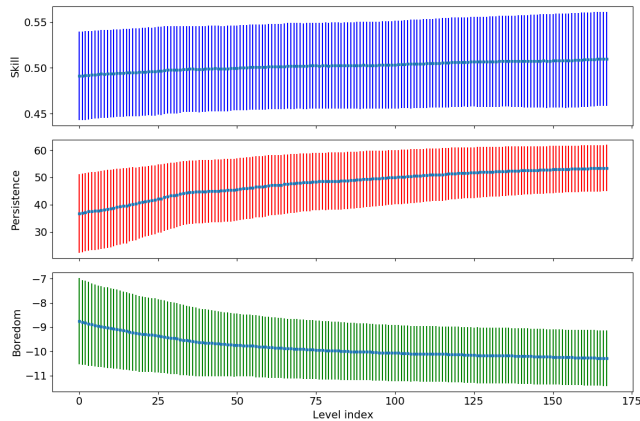


Figure 6: Evolution of the simulated player population’s mean and standard deviation of skill, persistence, and boredom attributes over game levels.

4.3.2 Design Rationale. We chose to model skill, persistence, and boredom based on our human player data and the literature on game engagement and player psychology. Figure 1 shows that pass rate and churn are related, with players being more likely to churn with low pass rate, in line with [24, 34]. A psychological explanation for this can be found in Self-Determination Theory [37], which posits that feelings of competence and mastery support intrinsic

motivation. Thus, repeatedly failing a level is likely to decrease the motivation to play. In AB Dream Blast, level pass rate depends on skill but with some randomness due to randomly spawned game objects, which motivates our stochastic skill model. The persistence coefficient is based on some players being more persistent and willing to keep trying when facing hard challenges [28]. Finally, game engagement tends to decline over time [44, 45]. One psychological explanation for this is that the novelty of a game fades over time, which makes players less likely to feel curious interest [42]. Considering boredom as the counterpart of curiosity, we model each simulated player’s tendency to get bored as probability of churning after passing a level. This makes the player increasingly likely to churn when progressing through multiple levels and the boredom random draw is repeated.

4.3.3 Optimizing Simulation Parameters. To fit the model to observed pass and churn rates from real players, we optimize three types of parameters: 1) the baseline pass rate predictor’s regression weights, 2) the means and standard deviations for the initial player population’s skill, persistence, and boredom attributes, and 3) the $\alpha, \beta, \theta, \gamma$ parameters of Algorithm 1.

Instead of jointly optimizing all parameters, we simplify by first independently fitting the baseline pass rate predictor parameters. The level difficulty d in Algorithm 1 line 11 is computed as $d = \text{normalize}(-\rho_b)$, where ρ_b is the baseline pass rate prediction, and the normalization is over all levels, transforming the difficulties between zero and one.

The rest of the parameters (10 in total) are optimized using CMA-ES [12], a derivative-free global optimization method, to minimize the following objective function:

$$f = \text{MSE}(\rho_p) + w_{\text{churn}} \text{MSE}(\rho_c), \quad (1)$$

where ρ_p and ρ_c are the pass rate and churn rate predictions of Algorithm 1, and MSE denotes computing the mean squared error over all levels. Since the churn rate varies in a much smaller range than pass rate, we employ the w_{churn} parameter to amplify the relative importance of predicting churn correctly. We set w_{churn} equal to the variance of human pass rates divided by the variance of human churn rates.

We use CMA-ES population size of 120 and optimize until there is no improvement for 100 iterations, which takes approximately two hours on an Intel core i7 2.11 GHz CPU and consumes about 170 Mb RAM.

5 EVALUATION

We use 5-fold cross-validation to compute prediction errors and compare the two models. More specifically, we always simulate over all levels, but leave out one fifth in computing the optimization objective function. Thus, human data from levels used for measuring validation error does not inform the model parameter optimization. This matches the intended use case where the model is optimized based on some existing level data, and a designer is using the model to predict pass and churn rates of new levels.

Table 1 shows the validation mean squared error (MSE) and mean absolute error (MAE) of the two models. The extended model improves churn prediction and performs approximately similarly in pass rate prediction. Note that the two models are complementary

Algorithm 1 Pass and churn simulation, run for each level

Input: level difficulty d , population P , standard deviations α, θ, β , learning rate γ

Output: pass rate ρ_p , churn rate ρ_c , evolved population P

```

1:  $\rho_p \leftarrow 0$                                  $\triangleright$  Initialization
2:  $\rho_c \leftarrow 0$                                  $\triangleright$  Initialization
3:  $M \leftarrow P.size$                              $\triangleright$  Remember population size
4: for player  $p$  in population  $P$  do
5:   draw  $s \sim N(\mu = p.skill, \sigma = \alpha)$ 
6:   draw  $t \sim N(\mu = p.persistence, \sigma = \beta)$ 
7:    $passed, churned = false$ 
8:    $nAttempts = 0$ 
9:   while not passed and not churned do
10:     $nAttempts \leftarrow nAttempts + 1$ 
11:    if  $s \geq \text{level difficulty } d$  then
12:       $passed \leftarrow true$ 
13:       $\rho_p \leftarrow \rho_p + 1/nAttempts/M$ 
14:      draw  $b \sim N(\mu = 0, \sigma = \theta)$ 
15:      if  $b < p.boredom$  then
16:         $churned \leftarrow true$ 
17:         $\rho_c \leftarrow \rho_c + 1/M$ 
18:         $P.remove(p)$ 
19:      end if
20:    else
21:       $s \leftarrow s + \gamma$                                  $\triangleright$  Learning
22:      if  $nAttempts > t$  then
23:         $churned \leftarrow true$ 
24:         $\rho_c \leftarrow \rho_c + 1/M$ 
25:         $P.remove(p)$ 
26:      end if
27:    end if
28:  end while
29: end for
30:  $P.add(\text{randomly select } M - P.size \text{ players from } P)$ 

```

instead of mutually exclusive: As the baseline pass rate predictor gives a slightly lower mean error, one might use its predictions as such, and only utilize the churn prediction of the extended model. The means and standard deviations are computed over the cross-validation folds. As CMA-ES results vary with random seed, we ran the cross-validation 5 times for the extended model, i.e., the mean and standard deviation are over a total of 25 optimization runs.

As shown in the scatter plots in Figures 7 and 8, the extended model is better able to discriminate between the early and later game levels. As players with low skill and persistence churn earlier, same as players who get bored easier, the average churn rate in early levels is higher than in the later levels.

To test the generalization to new levels added to the end of the game, we did an extra validation test that leaves out the last fifth of the levels during training. The pass and churn rate MSEs of these levels are 0.01003 and 0.00003, which are lower than the baseline 0.01073 and 0.00016. Naturally, our model can become less reliable when more new levels are added, especially if the levels feature novel gameplay. Thus, when releasing new levels, a game

developer should collect human player data and refit the population simulation parameters.

When carefully comparing the predicted and ground truth data in Figure 8, one notes that although our model captures the overall relation, some inaccuracy remains especially at low pass rates. We hypothesized that this is at least to some degree due to inaccuracies of the pass rates predicted by the AI gameplay. We tested this hypothesis by also running the simulation using $d = 1 - \rho_{\text{human}}$ as the difficulty estimate, where ρ_{human} is the human pass rate. This decreases churn prediction MSE by 71%. Figure 9 also shows how the scatter plots are visually more accurate. While the results indicate our AI game playing agents are not yet human-like enough, the results also suggest that our pass and churn simulation of the extended model is surprisingly accurate.

We performed an ablation study to investigate the effect of each simulation feature on the prediction results. Table 2 indicates that removing any of the features would make our churn model less accurate.

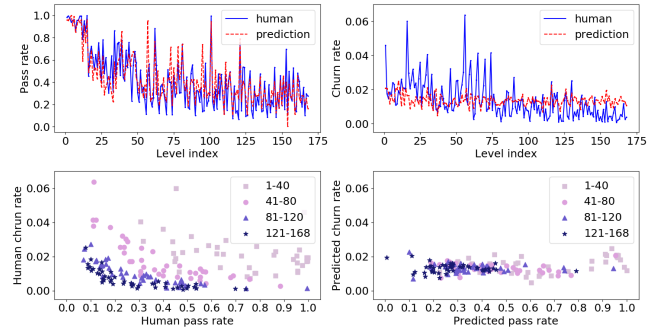


Figure 7: Prediction results of the baseline model. Scatter plot colors correspond to game level indices.

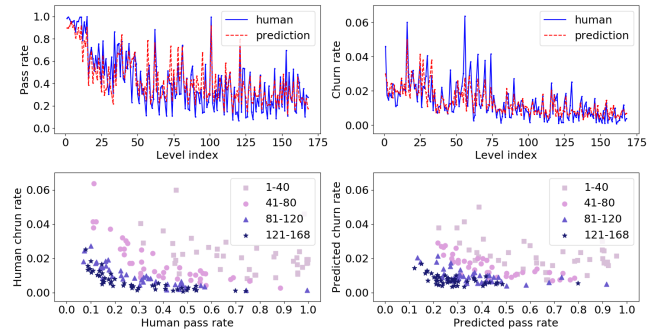


Figure 8: Prediction results of our extended model. Scatter plot colors correspond to game level indices.

6 LIMITATIONS AND FUTURE WORK

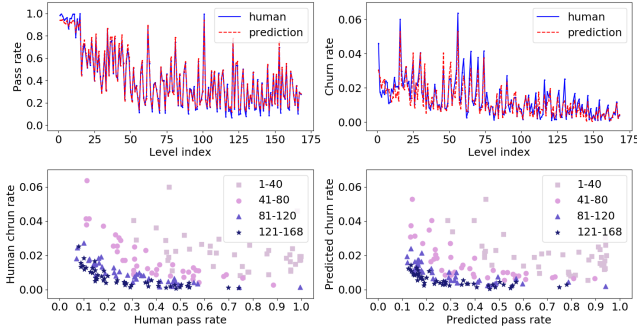
As a limitation, our model has not yet been validated in actual game design work, e.g., in testing new levels and identifying problematic levels before they are deployed to real players. This should be possible in future work, as our Deep Reinforcement Learning

Table 1: Mean squared errors and mean absolute errors of average pass rate and churn rate prediction.

Method	Validation MSE		Validation MAE	
	Pass rate	Churn rate	Pass rate	Churn rate
Baseline	$\mu = 0.02244$ $\sigma = 0.00803$	$\mu = 0.00013$ $\sigma = 0.00003$	$\mu = 0.11228$ $\sigma = 0.01663$	$\mu = 0.00866$ $\sigma = 0.00076$
Extended model	$\mu = 0.02320$ $\sigma = 0.00831$	$\mu = 0.00008$ $\sigma = 0.00002$	$\mu = 0.11467$ $\sigma = 0.01647$	$\mu = 0.00607$ $\sigma = 0.00073$

Table 2: The effect of different simulation components on the prediction results.

Method	Pass rate validation MSE	Churn rate validation MSE
All features	0.023	0.00008
No boredom	0.024	0.00009
No persistence	0.023	0.00010
No learning	0.032	0.00010
No random noise in skill and persistence	0.023	0.00014

**Figure 9: Extended model results when computing level difficulty based on actual human pass rates instead of those predicted by AI gameplay. Comparing this to Figure 8, one notes improved churn prediction for low pass rates. Scatter plot colors correspond to game level indices.**

game playing approach does not require any reference gameplay data and should thus generalize to new levels. Our cross-validation also suggests generalizability. As a downside, we noticed that the difficulty predictions of the DRL data are not yet entirely human-like especially for levels with very low pass rates. This is a topic to investigate in future work, especially as the accuracy of our population simulation and churn prediction does increase considerably when replacing DRL difficulty predictions with ones computed from ground truth human data.

Training a DRL agent on a new level is also slow, which imposes a bottleneck on simulated game testing efficiency. Ways to address this in future work include improving parallelism (e.g., [1]) and testing and comparing alternative game playing AI approaches, e.g., utilizing MCTS instead of DRL, or speeding up the DRL training process using transfer learning.

7 CONCLUSION

We have proposed a novel simulation model for predicting puzzle game level pass and churn rates. Our model combines both AI gameplay for level difficulty prediction, and simulating the evolution of a player population over game levels using simple computational models of skill, persistence, boredom, and learning. In terms of cross-validation error, the model outperforms a baseline model that predicts pass rates and churn based on AI gameplay data alone.

Our work shows how predictions produced by game playing Deep Reinforcement Learning agents can be enhanced by even a computationally very simple population-level simulation of individual player differences, without requiring retraining the agents or collecting new gameplay data for each simulated player. As training the DRL agents can be very slow, this is a foundational result that can empower future research and applications of player modeling and simulation-based game testing. Our model is also the first to make human-like pass and churn rate data emerge from AI gameplay in a way that captures how the relation of pass rate and churn evolves over a game’s level progression.

8 ACKNOWLEDGMENTS

We thank the anonymous reviewers for their thorough and constructive comments.

REFERENCES

- [1] Igor Adamski, Robert Adamski, Tomasz Grel, Adam Jędrych, Kamil Kaczmarek, and Henryk Michalewski. 2018. Distributed Deep Reinforcement Learning: Learn how to play Atari games in 21 minutes. In *International Conference on High Performance Computing*. Springer, 370–388.
- [2] Amazon. 2006. *Amazon web services*. <http://aws.amazon.com>
- [3] Sinan Ariyurek, Aysu Betin-Can, and Elif Surer. 2019. Automated Video Game Testing Using Synthetic and Human-Like Agents. *IEEE Transactions on Games* (2019).
- [4] Sinan Ariyurek, Aysu Betin-Can, and Elif Surer. 2020. Enhancing the Monte Carlo Tree Search Algorithm for Video Game Testing. *arXiv preprint arXiv:2003.07813* (2020).
- [5] Paul Bertens, Anna Guitart, and África Periañez. 2017. Games and big data: A scalable multi-dimensional churn prediction model. In *2017 IEEE Conference on Computational Intelligence and Games (CIG)*. IEEE, 33–36.
- [6] Valerio Bonometti, Charles Ringer, Mark Hall, Alex R Wade, and Anders Drachen. 2019. Modelling Early User-Game Interactions for Joint Estimation of Survival Time and Churn Probability. In *2019 IEEE Conference on Games (CoG)*. IEEE, 1–8.

- [7] Noshaba Cheema, Laura A. Frey-Law, Kourosh Naderi, Jaakko Lehtinen, Philipp Slusallek, and Perttu Hämäläinen. 2020. Predicting Mid-Air Interaction Movements and Fatigue Using Deep Reinforcement Learning. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, Article 572, 13 pages. <https://doi.org/10.1145/3313831.3376701>
- [8] Fernando de Mesentier Silva, Scott Lee, Julian Togelius, and Andy Nealen. 2017. AI-based playtesting of contemporary board games. In *Proceedings of the 12th International Conference on the Foundations of Digital Games*. 1–10.
- [9] Rovio Entertainment. 2018. *Angry Birds Dream Blast*. Game [Android, iOS].
- [10] Samuel J Gershman, Eric J Horvitz, and Joshua B Tenenbaum. 2015. Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science* 349, 6245 (2015), 273–278.
- [11] Stefan Freyr Gudmundsson, Philipp Eisen, Erik Poromaa, Alex Nodet, Sami Purmonen, Bartłomiej Kozakowski, Richard Meurling, and Lele Cao. 2018. Human-like playtesting with deep learning. In *2018 IEEE Conference on Computational Intelligence and Games (CIG)*. IEEE, 1–8.
- [12] Nikolaus Hansen. 2016. The CMA evolution strategy: A tutorial. *arXiv preprint arXiv:1604.00772* (2016).
- [13] Christoffer Holmgard, Michael Cerny Green, Antonios Liapis, and Julian Togelius. 2018. Automated playtesting with procedural personas with evolved heuristics. *IEEE Transactions on Games* (2018).
- [14] Aaron Isaksen, Daniel Gopstein, and Andrew Nealen. 2015. Exploring Game Space Using Survival Analysis. In *FDG*.
- [15] Arthur Juliani, Vincent-Pierre Berges, Esh Vckay, Yuan Gao, Hunter Henry, Marwan Mattar, and Danny Lange. 2018. Unity: A general platform for intelligent agents. *arXiv preprint arXiv:1809.02627* (2018).
- [16] Niels Justesen, Philip Bontrager, Julian Togelius, and Sebastian Risi. 2019. Deep learning for video game playing. *IEEE Transactions on Games* (2019).
- [17] Ildar Kamalidinov and Ilya Makarov. 2019. Deep reinforcement learning in match-3 game. In *2019 IEEE conference on games (CoG)*. IEEE, 1–4.
- [18] Todd B Kashdan and Paul J Silvia. 2009. Curiosity and interest: The benefits of thriving on novelty and challenge. *Oxford handbook of positive psychology 2* (2009), 367–374.
- [19] Seungwook Kim, Daeyoung Choi, Eunjung Lee, and Wonjong Rhee. 2017. Churn prediction of mobile and online casual games using play log data. *PLoS one* 12, 7 (2017).
- [20] Jeppe Theiss Kristensen and Paolo Burelli. 2019. Combining Sequential and Aggregated Data for Churn Prediction in Casual Freemium Games. In *2019 IEEE Conference on Games (CoG)*. IEEE, 1–8.
- [21] Guillaume Lample and Devendra Singh Chaplot. 2017. Playing FPS games with deep reinforcement learning. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- [22] Eunjo Lee, Yoonjae Jang, Du-Mim Yoon, JiHoon Jeon, Seong-il Yang, Sang-Kwang Lee, Dae-Wook Kim, Pei Pei Chen, Anna Guitart, Paul Bertens, et al. 2018. Game data mining competition on churn prediction and survival analysis using commercial game log data. *IEEE Transactions on Games* 11, 3 (2018), 215–226.
- [23] Conor Linehan, George Bellord, Ben Kirman, Zachary H Morford, and Bryan Roche. 2014. Learning curves: analysing pace and challenge in four successful puzzle games. In *Proceedings of the first ACM SIGCHI annual symposium on Computer-human interaction in play*. 181–190.
- [24] Derek Lomas, Kishan Patel, Jodi L Forlizzi, and Kenneth R Koedinger. 2013. Optimizing challenge in an educational game using large-scale design experiments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 89–98.
- [25] J Derek Lomas, Kenneth Koedinger, Nirmal Patel, Sharan Shodhan, Nikhil Poonwala, and Jodi L Forlizzi. 2017. Is difficulty overrated? The effects of choice, novelty and suspense on intrinsic motivation in educational games. In *Proceedings of the 2017 CHI conference on human factors in computing systems*. 1028–1039.
- [26] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [27] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529–533.
- [28] Joyce LD Neys, Jeroen Jansz, and Ed SH Tan. 2014. Exploring persistence in gaming: The role of self-determination and social identity. *Computers in Human Behavior* 37 (2014), 196–209.
- [29] Antti Oulasvirta. 2018. *Computational interaction*. Oxford University Press.
- [30] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. 2017. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 16–17.
- [31] Christopher Pedersen, Julian Togelius, and Georgios N Yannakakis. 2010. Modeling player experience for content creation. *IEEE Transactions on Computational Intelligence and AI in Games* 2, 1 (2010), 54–67.
- [32] África Perriñez, Alain Saas, Anna Guitart, and Colin Magne. 2016. Churn prediction in mobile social games: Towards a complete assessment using survival ensembles. In *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*. IEEE, 564–573.
- [33] Serge Petralito, Florian Brühlmann, Glena Iten, Elisa D Mekler, and Klaus Opwis. 2017. A good reason to die: how avatar death and high challenges enable positive experiences. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 5087–5097.
- [34] Erik Ragnar Poromaa. 2017. Crushing candy crush: predicting human success rate in a mobile game using Monte-Carlo tree search.
- [35] David Reguera, Pol Colomer-de Simón, Iván Encinas, Manel Sort, Jan Wedekind, and Marián Boguná. 2019. The Physics of Fun: Quantifying Human Engagement into Playful Activities. *arXiv preprint arXiv:1911.01864* (2019).
- [36] Shaghayegh Roohi, Jari Takatalo, Christian Guckelsberger, and Perttu Hämäläinen. 2018. Review of intrinsic motivation in simulation-based game testing. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [37] Richard M Ryan, C Scott Rigby, and Andrew Przybylski. 2006. The motivational pull of video games: A self-determination theory approach. *Motivation and emotion* 30, 4 (2006), 344–360.
- [38] Jesse Schell. 2008. *The Art of Game Design: A book of lenses*. CRC press.
- [39] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [40] Kun Shao, Zhentao Tang, Yuanheng Zhu, Nannan Li, and Dongbin Zhao. 2019. A Survey of Deep Reinforcement Learning in Video Games. *arXiv preprint arXiv:1912.10944* (2019).
- [41] Kennon M Sheldon, Andrew J Elliot, Youngmee Kim, and Tim Kasser. 2001. What is satisfying about satisfying events? Testing 10 candidate psychological needs. *Journal of personality and social psychology* 80, 2 (2001), 325.
- [42] Paul J Silvia. 2008. Interest—The curious emotion. *Current directions in psychological science* 17, 1 (2008), 57–60.
- [43] Jari Takatalo, Jukka Häkkinen, Jyrki Kaistinen, and Göte Nyman. 2010. Presence, involvement, and flow in digital games. In *Evaluating user experience in games*. Springer, 23–46.
- [44] Markus Viljanen, Antti Airola, Jukka Heikkonen, and Tapio Pahikkala. 2017. Playtime measurement with survival analysis. *IEEE Transactions on Games* 10, 2 (2017), 128–138.
- [45] Markus Viljanen, Antti Airola, Tapio Pahikkala, and Jukka Heikkonen. 2016. Modelling user retention in mobile games. In *2016 IEEE Conference on Computational Intelligence and Games (CIG)*. IEEE, 1–8.
- [46] Oriol Vinyals, Timo Ewalds, Sergey Bartunov, Petko Georgiev, Alexander Sasha Vezhnevets, Michelle Yeo, Alireza Makhzani, Heinrich Küttler, John Agapiou, Julian Schrittwieser, et al. 2017. Starcraft ii: A new challenge for reinforcement learning. *arXiv preprint arXiv:1708.04782* (2017).
- [47] Robert C Wilson, Amitai Shenhav, Mark Straccia, and Jonathan D Cohen. 2019. The Eighty Five Percent Rule for optimal learning. *Nature communications* 10, 1 (2019), 1–9.
- [48] Wanshan Yang, Ting Huang, Junlin Zeng, Gemeng Yang, Jintian Cai, Lijun Chen, Shivakant Mishra, and Youjian Eugene Liu. 2019. Mining Player In-game Time Spending Regularity for Churn Prediction in Free Online Games. In *2019 IEEE Conference on Games (CoG)*. IEEE, 1–8.
- [49] Georgios N Yannakakis, Pieter Spronck, Daniele Loiacono, and Elisabeth André. 2013. Player modeling. (2013).
- [50] Yunqi Zhao, Igor Borovikov, Fernando de Mesentier Silva, Ahmad Beirami, Jason Rupert, Caedmon Somers, Jesse Harder, John Kolen, Jervis Pinto, Reza Pourabolfghasem, et al. 2019. Winning Isn't Everything: Enhancing Game Development with Intelligent Agents. *arXiv preprint arXiv:1903.10545* (2019).