Simulation-driven gym layout optimisation

Andrej Jočić, Matic Stare, and Martin Starič

Collective behaviour course research seminar report

This report describes our work on simulation-driven gym layout optimisation. We begin with a brief overview of the problem, followed by a description of our agent-based model of gym-goer behaviour. Next, we describe a constrained multi-objective optimisation procedure for finding optimal gym layouts with genetic algorithms, and present some preliminary results. We conclude with a discussion of our solution's limitations and possible future work.

gym layout | crowd simulation | gym-goer behaviour | gym traffic

Introduction

In this report we describe our work on gym layout optimisation. A common occurrence in gyms during peak hours is over-crowding of popular equipment and lots of wandering around searching for a free machine. Through simulating gym-goer behavior, we aimed to identify the most effective arrangement of exercise equipment to improve the flow, accessibility, and overall customer satisfaction. Our goal is to offer practical insights for gym owners, managers, and designers seeking to optimise their facility layout for the benefit of their clients and business success.

To simplify the problem we will limit ourselves to simulating the behaviour of body-builders (performing resistance exercises for muscle hypertrophy). In this case, a typical workout routine partitions the body into several sets of muscles, cycling through them on a per-workout basis. To further simplify the problem, we will only simulate the *push-pull-legs* (PPL) partitioning[1], which is a popular choice for beginners and intermediate lifters. We will model the gym as a rectangular grid of cells with pre-set equipment locations, distributed similarly as isles in a grocery store.

Related work. The only publication related to gym layout optimisation we could find was ~~ref.~~ [1], which assumed all gym clients had fixed-order workout routines (including ones with weight loss as a goal) and optimised a circular gym layout to minimise backward movement. Unfortunately, this does not give a good foundation for our work, as the assumptions diverge too far from what we are trying to model.

Thus we started from a crowd modelling review [2] for basic model design principles. We also found two articles on incremental urban layout optimisation [3, 4] useful when designing the optimisation procedure, especially the cost function.

Methods

We used Python's *Mesa* library for simulation, analysis, and visualization of agent-based models.

Gym layout model. A gym is represented as a subclass of `mesa.Model` with two rectangular cell grids: the agent layer (where gym-goers move around) and the equipment layer (where machines are statically placed for one simulation cycle). One cell on the grid's border serves as the locker area entrance, where agents enter the gym. The agents can only move around in cells that are not occupied by a piece of equipment, but there can be multiple agents in a single cell. We do not feel that agent collision resolution is necessary for realistic modelling of gym traffic.

For generating a concrete placement of machines, we have to assign certain cells to be equipment locations (where one machine is placed) while making sure we don't create unreachable areas. We have developed some basic layout templates mimicking the arrangement of isles in a grocery store, with sensible walkways in between. In this case, layout optimisation comes down to choosing the best assignment of machines to equipment locations.

Gym-goer behaviour model. A gym-goer is represented as a subclass of `mesa.Agent` with a training checklist (multiset of muscles) and its current state. The agent's goal is to exhaust the checklist as quickly as possible and exit the gym.

---

[1]A *push* workout consists of pushing movements, mainly targeting the chest, frontal deltoids (shoulders), and triceps. Pull muscles are those involved in pulling movements, such as the back musculature, rear deltoids, and biceps. The content of a leg workout is self-evident.
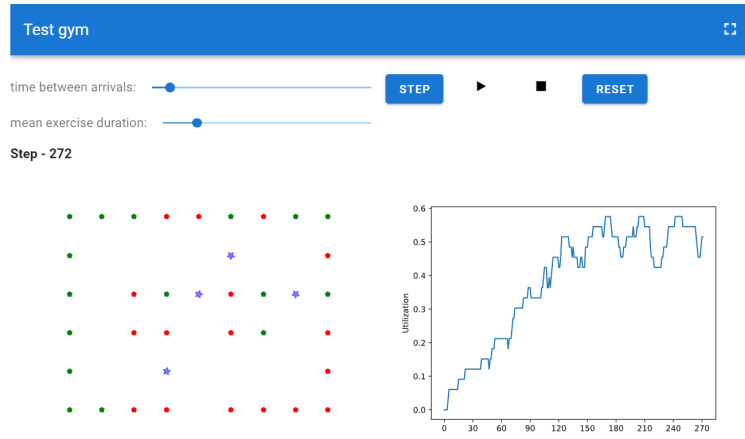
Figure 1. A gym traffic simulation cycle frozen at 272 time steps. On the left we can see the gym environment, generated from a layout template with the entrance at the bottom. Red pentagons are occupied machines, while green ones are free. Agents that are currently walking around in search of a machine are shown with blue stars. On the right, we can see the time progression of one of our gym evaluation metrics (described in a later section).

A simulation cycle begins with an empty gym. Then, one agent is spawned at the entrance every t_i time steps, where t_i is a fixed inter-arrival time given as a parameter to the simulation. When an agent is created, it is assigned a workout routine (e.g. push) and a corresponding workout plan (e.g. chest, chest, triceps, frontal deltoids). In our experiments we sampled the routine selection from a uniform probability distribution. The workout plans are constructed ad hoc, in keeping with general workout programming practice. The agent also keeps track of the machines it has already used, so it can avoid using the same machine twice in one workout.

Upon entering the gym, an agent starts in state *searching*, where it looks for a free piece of equipment it can use to check any muscle of its training checklist. A path to the nearest free machine is found using the A* algorithm, which the agent begins to follow. If the target machine is no longer free upon arrival, the agent continues in state *searching*. When a free machine is found, the agent enters state *exercising* and occupies the machine for a certain amount of time. In our experiments, we used a deterministic exercise duration of 30 time steps for all exercises. After the time has elapsed, the agent re-enters state *searching* or exits the gym if it has completed its workout.

Layout optimisation. We formalized our problem as a constrained multi-objective optimisation problem, where a layout is constrained by a given layout template. This constraint makes the optimisation problem more tractable, and we can also ensure that no unreachable areas are ever created. With multiple optimisation objectives, we didn't have to combine all aspects of a well-designed gym into one number and thus make arbitrary decisions about the relative importance of each aspect.

The procedure begins with a set of random layouts (meaning random assignments of machines to locations in the template), and evolves them with a genetic algorithm. We used Python's *PyGAD* library for genetic algorithms, since it supports multi-objective optimisation with the NSGA-II algorithm [5]. This converges to a set of Pareto-optimal solutions, meaning that no other solution can improve on any of the objectives without worsening at least one of the others.

A chromosome in our genetic algorithm is a list of exercise machines. To evaluate a chromosome, we first sequentially assign machines to locations in the layout template. The locations are ordered in a way that largely preserves the spatial structure of the template, so that a k-point crossover operation doesn't completely destroy the layout topology. For example, the 4-isle layout template (Figure 4b) is ordered as follows: first come the locations around the walls in counter-clockwise order, starting from the bottom-left. Then, all the isles are traversed in row-major order.

After a gym is instantiated, its fitness values are calculated by running a simulation cycle on it and extracting selected metrics (described in the next section). These metrics constitute the compound objective function. Note that random mutations and crossover can easily lead to layouts on which some of our hand-crafted workout routines cannot be completed; we simply set their fitness to negative infinity, effectively discarding them from the population. In this way, roughly 15% of solutions were discarded from each generation on average.

We used population sizes around 50 in our experiments. Since this optimisation problem has a complex fitness landscape, we would ideally use even larger ones, but

we were limited by the performance of our simulation. We used a relatively small inter-arrival time $t_i = 2$ to speed up the convergence of optimisation objectives, stopping each simulation after 300 steps. The random mutation rate was set to 15%, and we selected the 2-point crossover operator.

Objective functions. The layout quality metrics which we maximized were:

- agent *efficiency*: proportion of time spent exercising, averaged over all agents in the simulation cycle,

- gym *utilization*: proportion of machines in use at any given time, averaged over the whole cycle, and

- gym *crowdedness*: proportion of pairs of agents that are "too close" to each other at any given time. This was defined as having an inter-cell manhattan distance ≤ 2, and computed only on agents in state *searching*. After being averaged over the simulation cycle, this value also had to be inverted since PyGAD only supports fitness maximization.

Results

We tested our optimiser on a few small layout templates. First we used a template which has machines around the walls and two 2-by-2 islands of machines in the middle (Figure 2b). The fitness functions converged at around 40 generations. Note that the utilization criterion (2a, blue line) decreases as the other two increase. This represents the trade-off between customer satisfaction and gym revenue, as an efficiently utilized gym is cheaper to set up and maintain, but ends up being more crowded.

**(a)** Evolution of fitness functions over generations.

(b) A Pareto-optimal gym layout.

Figure 2. Optimisation constrained to a 2-island layout template. Machines (on the right) are labeled with name abbreviations, see the `Equipment` class in `gym_model.py` on our GitHub repository [6] for the list of machines we used. See Figure 3 for an interpretation of the colors.

Figure 3. The colormap we used for visualizing a gym layout. Each machine is colored by the muscle it trains, where muscles are mapped onto the color spectrum in (roughly) the same order as they appear in the human body (from fingertips to heels).

Next, we ran the optimiser on a template with 4 islands of machines (Figure 4b). Using the same parameters as before, the optimiser converged after around 110 generations. Due to increasing gym size and a more complex search space, the performance of our simulation prevented us from experimenting with larger layouts.

(a) Evolution of fitness functions over generations.

(b) A Pareto-optimal gym layout.

Figure 4. Optimisation constrained to a 4-island layout template.

Discussion

We successfully implemented a gym layout optimisation framework, but our results stand only as a proof of concept. We made many simplifying assumptions which likely render our model too unrealistic to be of practical use. Since we didn't invest much time into code optimisation, we couldn't perform comprehensive experiments.

Future work. Our gym layout model could be made more realistic by dropping the constraints on layout topology and adding support for shared resources (e.g. free weights) and multipurpose equipment. We could also implement more fine-grained agent behaviour, such as parallel use of gym equipment (where one agent is working while the other is resting), random inter-arrival times (e.g. according to a Poisson process) and random exercise durations, as well as more realistic routine selection (perhaps with parameters estimated from real-world data).

It would be interesting to let our optimiser come up with a layout from scratch, without the need for hand-crafted templates. Perhaps this could be done with a Metropolis-Hastings state search variation, which has been successfully applied to urban layout optimisation [3, 4]. Another interesting idea ~~in ref. [3]~~ is the use of nonlinear regressors to predict the quality of a layout based on topological features, circumventing the need for computationally expensive simulations. We could also try to account for the cost of equipment (given by a gym designer) in the optimisation criteria.

CONTRIBUTIONS. Everyone worked on the report and gym-goer behaviour model. Andrej Jočić worked on model implementation and the layout optimiser. Martin Starič implemented the pathfinding algorithm. Matic Stare made the project presentation.

Bibliography

1. Küçükdeniz T (2022) Gym layout modeling and optimization. *Management Science Letters* 12(3):185–192.
2. Yang S, Li T, Gong X, Peng B, Hu J (2020) A review on crowd simulation and modeling. *Graphical Models* 111:101081.
3. Feng T, Yu LF, Yeung SK, Yin K, Zhou K (2016) Crowd-driven mid-scale layout design. *ACM Trans. Graph.* 35(4):132–1.
4. Mathew CT, Knob PR, Musse SR, Aliaga DG (2019) Urban walkability design using virtual population simulation in *Computer Graphics Forum.* (Wiley Online Library), Vol. 38, pp. 455–469.
5. Deb K, Pratap A, Agarwal S, Meyarivan T (2002) A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation* 6(2):182–197.
6. (2024) Github repository: gymulator (https://github.com/andrejjocic/gymulator).