

Behaviour Tree Evolution by Genetic Programming

– Learning Novel Bot Behaviours in a 2D Top-Down Arena Shooter –

Martin Estgren <mares480@student.liu.se>
Erik S. V. Jansson <erija578@student.liu.se>

Linköping University, 12th November 2017

Abstract

Behaviour trees are a popular model for representing the decision-making and plan execution process for NPCs in video games. These are built by hand, and often require expertise to craft into interesting and intelligent behaviours. These types of behaviour trees usually don't adapt well to new environments, and need many tailor-made trees for each situation.

In this paper we demonstrate how to generate the trees by using genetic programming; allowing us to essentially evolve novel behaviours automatically in our testbed. Instead of specifying low-level actions and conditions, we use high-level definitions. This leads to faster fitness convergence, and allows us to skip the bloat control and tree pruning steps used by other similar behaviour tree evolution methods. Results show that the evolved trees reliably defeats our hand-crafted behaviours by the 25th generation.

Contents

1	Introduction	
1.1	Proposed Approach	
2	Background Theory	
2.1	Behaviour Trees	
2.2	Path Finding with A*	
2.3	Genetic Programming	
3	Implementation Details	
3.1	Game Architecture	
3.2	Behaviour Trees	
3.3	Genetic Programming	
4	Results and Screenshots	
4.1	Generated Behaviours	
4.2	Behaviour Fitness	
4.3	Survival Ratio	
5	Discussion and Outlook	

1 Introduction

In interactive media such as video games, there is often a need for simulating seemingly complex agent behaviours in (soft) real-time. In a modern fast-paced computer game for example, the render, physics and artificial intelligence updates all need to complete in a total of less than 16ms to provide a half-decent experience. These resource constraint have spawned some very clever techniques to enable complex behaviours for autonomous agents while minimizing their required computational time.

These techniques often require the bot behaviours to be manually defined, and in some cases, are quite hard to extend when new behaviours are needed later in the game development process. This has been shown to be true by industry researchers, such as *Dawe et al* [4], regarding the *finite state machine architecture* (which we talk more about later).

In this project we have explored how one of the more widely used techniques, *behaviour trees*, can be extended to allow for not only hand-crafted complex behaviours, but for organic behaviours, tailored to the specific game domain by using a *reinforcement learning* approach. More specifically, we evolve the behaviour trees by using *genetic programming*. It's by no means a novel approach, but there still has been relatively little research in the area. In our *behaviour evolution algorithm*, based on the work of *Ögren et al.* [3], we instead use *high-level behaviours*, which removes the need for *bloat control* or *pruning*. Leading to a simpler algorithm than *Ögren et al.* [3].

Alongside the proposed method in Section 1.1, a testbed consisting of a top-down arena shooter was built. It features a sufficiently complex environment to support interesting behaviours by the enemy bot. We describe the features of our testbed in Section 4.

This report starts off in Section 2 by giving a brief overview of the relevant techniques and related work necessary to follow our reasoning and findings. In Section 3 we describe the details needed to implement these techniques in practice, along with the testbed architecture we’ve used. This gives the reader information on any peculiarities and pit-falls. We follow this by showing in Section 4 the testbed, and give some generated offspring generated by our method along with measurements of its efficiency. Finally, we reflect on our method, giving downsides, and present some possible future work in Section 5.

1.1 Proposed Approach

Apply *genetic programming* to the *behaviour trees* by creating *high-level actions and conditions* nodes. Ideally this should allow for faster convergence rates and somewhat shallower generated behaviour trees, eliminating the need for *bloat control* and *pruning*. Since BTs are trees, the *mutation mechanism* should modify *control-flow structure* and *leaf-node settings*, while the *crossover mechanism* will swap sub-trees.

2 Background Theory

For an autonomous agent to give the illusion of being intelligent, many different systems need to interact with each other. In this chapter we describe the theory behind our three primarily used techniques: *behaviour trees*, *path finding*, *genetic programming*. There are entire books written on these techniques, hence we’ll only describe each of these superficially, if you want a deeper treatment, see our references.

2.1 Behaviour Trees

Defining an accurate *automatic planner* model is often impractical and usually overkill for real-world applications. Especially in games, where very simple models are enough to give the illusion of intelligence. In games, the most well-known *behaviour selection algorithms* are *FSMs*, *GOAPs*, *HTNs* and *BTs* [4].

Finite State Machines (FSMs) are simple but hard to extend with additional states, resulting in an exponential increase of transitions [4]. While both *GOAP (Goal-Oriented Action Planner)* and *HTN (Hierarchical Task Network)* are powerful models,

these don’t allow agents to explore new task definitions, and “only” provide an agent with novel ways of solving an already defined set of tasks. Planning has seen some use in the games industry. Most notably in the game *F.E.A.R* [11] where adversarial agents used a planning system to combat the player.

Behaviour Trees (BTs) on the other hand, when comparing to *FSMs*, provide advantages in terms of modularity, reactivity and scalability. And more importantly, since they are a type of tree, allow for integration with the genetic programming approach. They were first popularized in *Halo 2* [7], but are now applied in other area as well, such as robotics.

Behaviour Tree Evolution (BTE) is a technique which uses some form of genetic programming on BTs. It’s a relatively new technique that has been applied to *strategy games* and *platformers*. Some of the previous work is due to *Lim et al.* [8], *Colledanchise et al.* [3] and a recent thesis by *Hoff et al.* [6]. To our knowledge, our paper is the first to attempt to apply BTE in a *shooter game*, and the first to use *high-level behaviours* for a simpler algorithm.

In the *Behaviour Tree Starter Kit* [2] and in *Chris Simpson’s Gamasutra article* [13], the BT is a tree-like data structure which describes the decision-making process of the agent. It’s evaluated pre-order from the root of the tree, in each logic update tick.

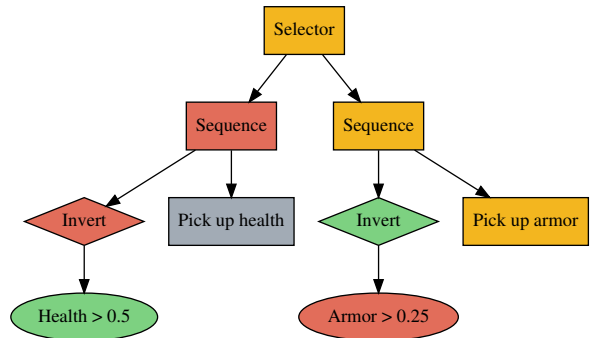


Figure 1: Example of Hand-Crafted Behaviour Tree

In the coming page you’ll find the most common types of *concrete nodes* in a behaviour tree, in particular those found in the *BTSK* [2]. Only *action* and *condition nodes* are *leaf-nodes*, and are *concrete nodes*. All others are either *composite* or *decorators (abstract nodes)*, which need to be instantiated to become concrete, such as the *sequence* or *inverter*; that have one/more children (it’s never a leaf node).

- **Selector:** executes children from left \rightarrow right until *one child* is found to be either *running* or has already *succeeded*, and returns this status. Otherwise, if none matches, it returns a *failure*.
- **Sequence:** executes children from left \rightarrow right until *one child* is found to be either *running* or *failed*, and returns this status. If *all children* have *succeeded*, then set status to *success*.
- **Inverter:** executes the child node and *inverts the resulting state* when it has finished running.
- **Succeder:** when child has finished running, returns *success* regardless of the child’s status.
- **Condition:** a leaf-node that queries the world state, if the query returns true, we give *success*, if it’s false, we instead give *failure*. This node is usually instantaneous, and can’t be *running*.
- **Action:** a leaf-node which modifies the world state, and returns *success* if the action had the intended side-effect, *failure* if couldn’t complete or failed the action and *running* if not done yet.

We find that it’s easier to understand behaviour trees with an example. We’ll be using the tree in Figure 1 as an example (it was produced within our testbed using a Graphviz library for Java). The yellow nodes means the node state is *running*, the red nodes have *failed*, and green nodes have *succeeded*.

At a high-level, this BT checks the agent’s health and armour, and if it’s below a threshold, it will go and pick up the crates which replenish these stats.

It does this by starting at the root node, *selector*, and executing the left child, the *sequence*. It will execute its left child, which is the *inverter*. At last we reach a leaf-node, the *health condition node*. This will probe the world for the state of the agent, and return *success* because we’ve found that it’s above 50%. This result will propagate upward to its parent, the inverter, which will set its state to *failed* (inverting the result of the child). We’re back at the left sequence, and because one of the children have failed (the inverter), the sequence has *failed* as well, and won’t run the right child (gray means the node has never been evaluated). Finally, because the selector’s left child has *failed*, we execute the right branch too (remember, at least one child must succeed!). We use the same steps as before, but get stuck in “pick up armour” until that action has finished *running* in the world (has succeeded or failed).

2.2 Path Finding with A*

Agent path finding is done using the A^* graph traversal algorithm proposed by Peter E. Hart, Nils J. Nilsson, and Bertram Raphael [5] as an extension to *Dijkstra’s algorithm* which, given an *admissible heuristic* and non-negative costs, will find the path from a node n_0 to a node n_g , with the lowest cost.

Calculating the expected cost of a path $n_0 \rightarrow n_k \rightarrow n_g$, from n_0 to n_g via n_k can be done with the following path finding cost function:

$$f(n_k) = g(n_k) + h(n_k)$$

where $g(n_k)$ is the true cost from n_0 to n_k and $h(n_k)$ is a heuristic approximation of the cost from n_k to n_g . This allows us to explore the potentially most optimal paths first before considering the others.

In this project, the *Euclidean distance* is used as the heuristic $h(n_k)$. This, since the path finding is performed in a search-space where each node is defined as a point in the testbed’s 2-D world-space.

Additionally, compared to regular A^* , the path-finder utilizes an *influence map* to help the agent avoid “suicidal paths”, even if the selected path is optimal given the heuristic function. Our *influence map*, like the ones shown in *Dave Mark* [9], take into account the visibility of each node in relation to the adversarial game agents’ line-of-sight, and therefore updates the expected cost as:

$$f(n_k) = g(n_k) + h(n_k) + i(n_k)$$

where $i(n_k)$ is the cost introduced by the influence map. This gives the the result that nodes within the adversarial agent’s line-of-sight have a much higher associated traversal cost, and should be avoided.




Figure 2: optimal path Figure 3: tactical path

Above are two example figures: Figure 2 shows an optimal path using regular A^* , and the other, Figure 3 shows the optimal path when taking into account the influence map; sneaking behind cover.

2.3 Genetic Programming

In many scenarios you try to optimize a function: $f(x_0, x_1, \dots, x_n)$, where some or all of the inputs are of a discrete nature (ordinal or nominal values). These types of problems are often hard to optimize using techniques developed for continuous variables, such as *gradient ascent/descent*, e.g. in a *perceptron*.

Genetic programming is a technique that was introduced by *Barricelli et al* [1], as a way to optimize computer programs by encoding the parameters to be optimized as *genetic representations* which could be processed by traditional *evolutionary algorithms*.




Figure 4: The *genetic representation* is often done using a tree-like structure in *genetic programming*. Image is taken from: Wikipedia (public domain).

The most general form of the *evolutionary algorithm* can be described in the following steps:

1. Generate an initial population by creating a initial pool \mathcal{P}_0 , of individuals with randomized genetic representations (context-specific data).
2. Evaluate the *fitness* for each individual $x_i \in \mathcal{P}_t$. This usually requires a domain-specific *fitness function*: $f(x_i) \rightarrow f_i \in \mathbf{R}$, which will be described in further details later in this chapter.
3. Remove individuals with poor *low fitness score* and *regenerate* population pool: $\mathcal{P}_{t+1} = \Phi(\mathcal{P}_t)$.

The fitness function for this project is composed of a *linear combination* of desirable behaviours collected during the *fitness evaluation* simulation, such as, amount of *damage dealt* to the adversary, whether the agent with a specific BT *survived*, and how many *weapons the agent has picked-up*. In the implementation section we'll specify f_i more concretely, for now we just explain the general theory behind genetic programming and how to find \mathcal{P}_{t+1} .

Regeneration of the population pool for the next generation, \mathcal{P}_{t+1} , is done using the three functions:

Selection $\Phi_s(\mathcal{P}_t) \rightarrow x_i \in \mathcal{P}_t$ individuals with high *fitness score* in the new population \mathcal{P}_{t+1} . In this project, *selection* is done proportionally to the *fitness score* for each individual. Each individual x_i has a probability $P(x_i|f_i)$ to be selected for the next generation \mathcal{P}_{t+1} according to some random process.

$$P(x_i|f_i) = \frac{f_i}{\sum_{j=1}^N f_j} \text{ where } N = |\mathcal{P}_t|$$

Crossover $\Phi_c(x_i, x_j) \rightarrow \hat{x}_i, \hat{x}_j$ individuals from the new population which have been crossed over, in the case of *genetic programming*, *crossover* is done by swapping random *sub-trees* in each BT. All sub-trees have an equal probability of being selected to crossover, creating new individuals in \mathcal{P}_{t+1} .

Mutation $\Phi_m(x_i) \rightarrow \hat{x}_i$ a subset of the new population, but slightly tweaked with a random parameter or by adding/removing random child nodes. If no mutation is done, then the genetic programming might converge to some local optima because of monoculture, and never reach a global optimum.

Once these steps have been performed, a new generation of individuals have been produced, which can then be used to repeat the optimization process until we converge to a population with behaviours close the global maximum of the fitness function f_i .

3 Implementation Details

After presenting these theoretical methods we'll now describe the additional changes necessary to make these work in practice (as per our implementation).

Both the top-down shooter game (target *testbed*) and our *behaviour tree evolution* (the *technique*) are written in *Java* and use *libGDX* for graphics/audio. This has allowed us to quickly produce a prototype, and while the game itself isn't the main topic of this paper, a rough outline can be found in Section 3.1.

Every *entity* in the game can be transformed into an intelligent agent with a `BotInputComponent`, which associates a given *behaviour tree* with a entity. It hooks together with the *path finder* to enable the entity to *traverse* & *interact* with the environment. With GP, additional BTs can be generated and used.

3.1 Game Architecture

Given the very limited implementation time of this project, the architecture of the game was designed to allow for fast prototyping iterations, and isn't necessarily based on strong and scalable software design philosophies. The structural design is based on the *entity component system* design, which has seen high use in the contemporary games industry and is being used in high-profile projects such as the *Unity* game engine. The primary reference implementation is taken from the book *Game Programming Patterns* written by *R. Nystrom* [10] remixed with the component system from *M. Boströms' Speed Coding Zelda*¹. The resulting implementation can be seen in the source code in the package: `se.sciion.quake2d.level` and its sub-folders.

An *entity component system (E-C-S)*, as the name might suggest, consists of three interacting class types. The *entity*, the *component*, and the *system*.

The *entity* is defined as a class consisting of a series of *components*, together with functions for querying information about the entities' components. This allows the entities' components to query their parent (the entity itself) regarding other components within this entity. Meaning, components that heavily depend on each other have high coupling, while independent components don't have any coupling. This gives us a very modular design, and gives the expressive power to easily add new components to an entity on-the-fly. There are several other ways to solve the problem of system-system communication, like message-passing, but we've decided to take the "easy way out" and just pass references around to the relevant objects.

In Figure 5 are the three different types of entities in our game: *user-controlled players*, *AI-controlled agents* and *health/armour/weapon crate pickups*. As can be seen, adding functionality to each entity is just a matter of adding/swapping a component. The only difference between a *player* and an *agent* is the method of input, *UserInput* vs *BotInput*, which in our opinion perfectly shows the benefits of a E-C-S.

The *BotInputComponent* has a behaviour tree which is updated with `tick`, which executes the motions specified in the *current target path* (basically a stack of locations to visit to reach the *target location*) which were decided by the behaviour tree.

¹<https://github.com/MilleBo/SpeedCodingZelda>

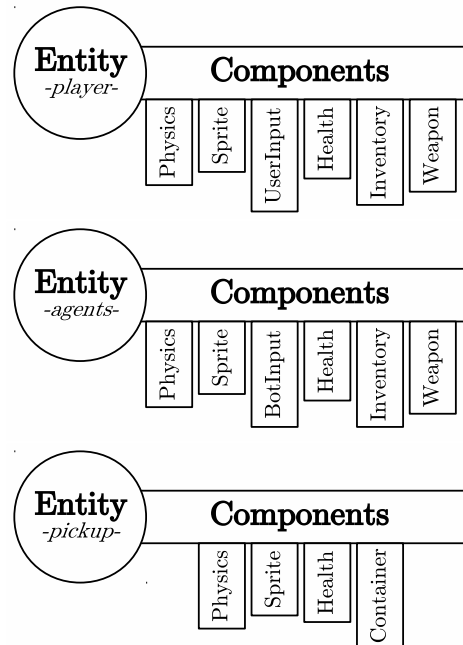


Figure 5: List of Components in Entity

Outside the entity, a set of *systems* are running, which manage inter-entity communication such as *path-finding*, *physics*, and *gameplay*. As we touched upon before, these systems are passed as references to the components that are interested in them, allowing for the components to decide what to do with the data provided by the systems without caring about the system's implementation details.

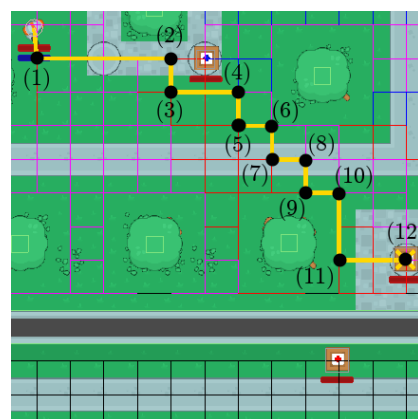
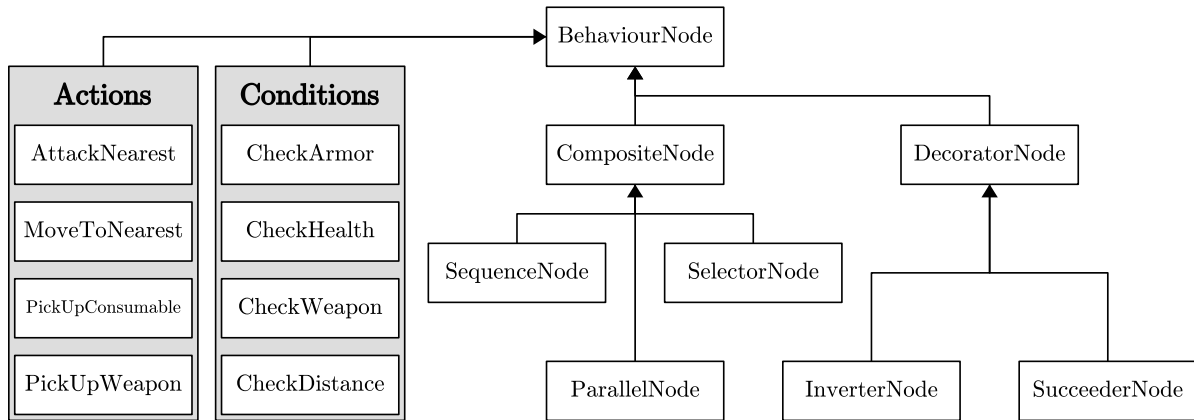


Figure 6: Visualization of the path finding stack \mathcal{F} . If there are no obstacles in the way (line-of-sight check), the agent takes a straight-line path instead.



3.2 Behaviour Trees

Inside of `se.sciion.quake2d.ai.behaviour` you'll find the implementation of our BTs. They are specified using inheritance, as in the figure above. In our variant, only the leaves above are non-abstract.

All are derived from the `BehaviourNode`, which provides the `State`, and methods for updating states, by evaluating it and its children (if it has any). This is done with the `tick` function, which is called for each agent's behaviour tree and once for each frame.

Algorithm 1 Pseudo-Code for the BT's Tick Step

```

Require: some initial State to the behaviour tree
if State  $\neq$  Running then
  doBehaviourEnter()
end if
State  $\leftarrow$  doBehaviourUpdate()
if State  $\neq$  Running then
  doBehaviourExit()
end if

```

Each of these entities' behaviour tree are built from the bottom-up and linked. There is only one root node which needs to be evaluated, since it will recursively evaluate all it's children when requested.

The standard non-leaf nodes are implemented in a straight-forward way following Section 2.1. Instead, we continue by explaining the *actions* & *conditions* we have implemented, and how these communicate with the other systems to be able to query the world.

We've adopted a fairly simple solution, where the leaf-node is given a reference to systems that it

wants to communicate with, and the entity owner.

For example, if a BT wants a *AttackNearest* node beneath a *Succeder*, we'd first have to hand over a reference to the *path finder* and a *target tag* to the node's constructor, such as: *AttackNearest(path-finder, tag)*. Only after that are we allowed to execute: *SuccederNode(attack-nearest)*, which we can utilize to construct the next parent layer above it.

Below is a list of all actions and conditions we've implemented in our game, along with a description:

- **AttackNearest(*tag*):** finds the closest entity with *tag* and *path find* to it. Shoot with weapon. *Success* if we were able to fire, *false* otherwise.
- **MoveToNearest(*tag*):** finds the close entity with *tag* and *path find* to it. Stops when close. *Success* when we're close, *fails* if no valid path.
- **PickUpConsumable(*type*):** finds the closest consumable of *type* (*health*, *armour* and *boost*), *path find* to it and *pick it up*. *Fails* if it couldn't.
- **PickUpWeapon(*type*):** finds a close weapon of *type* (*shotgun*, *rifle* & *sniper*). *Path find* to it & try to *pick it up*. *Succeeds* when picked up.
- **CheckStatus(*ratio*, *type*):** checks the entity owner's *status* (*health* or *armour*), and returns *Success* if above the *ratio*, or *fail* if its below.
- **CheckWeapon(*type*):** checks entity owner's *inventory* for weapon *type*. *Success* if we got it.
- **CheckDistance(*tag*, *r*)** checks if entity with *tag* is above *r* units away. *Succeeds* query if so.

3.3 Genetic Programming

Since the genetic programming has to be done on the trees during runtime, we need a way to dynamically generate trees with varying structure (changing composite nodes) and behaviours (changing leaf nodes). The solution used in this project is to create a set of *prototype behaviour nodes* \mathcal{S} , one for each node type. This set can then be used to create the trees by building them from the root. The prototypes are implemented using the *prototype design pattern* as presented by *R. Nystrom* [10], meaning, we have a set of correctly pre-specified instances which are then cloned & modified to suit our needs.

Tree Generation

Each tree starts by first generating a random node $n \in \mathcal{S}$ as the *tree root*, and if the node is a *composite*, continue recursively until a complete tree is formed.

All nodes have a *randomization function* which for each node type randomizes it in different ways. *Leaf nodes* have a trivial randomization function which only tweaks parameters such as *thresholds* (e.g. 0.5 or 0.2) and *game tags* (e.g. *player* or *crate*).

$$r(n) \rightarrow \hat{n}$$

is the randomization function for leaf node n to \hat{n} .

Composite nodes $\mathcal{C} = \{n_0, \dots, n_i\}$ are *randomized recursively*. First we randomize the set of children: $\mathcal{C}_r = \{n_0, \dots, n_m\}, n \in \mathcal{C}$ and $\mathcal{C}_a = \{n_0, \dots, n_j\}, n \in \mathcal{S}$. \mathcal{C} is then redefined as $\mathcal{C} = (\mathcal{C} \cup \mathcal{C}_a) / \mathcal{C}_r$, adding the set \mathcal{C}_a and removing the nodes in the set \mathcal{C}_r . Once this is done, each child nodes' respective *randomize function* is called:

$$r(\mathcal{C}) \rightarrow \forall_x r(n), x \in \mathcal{C}$$

is a randomization for the composites' children n .

Tree Mutation

As with *tree randomization*, the *mutation of trees* is done in a recursive manner, where each node mutates itself, and in the case of composite nodes its children. The mutation only differs from the tree randomization in that nodes are only slightly changed with a certain probability, whereas the randomization can completely change all nodes in a tree. In Figure 7 you can see *leaf-node mutation* in action, and in Figure 8 the *mutation of composites*.




Figure 7: Example of Mutation for the Leaf Node




Figure 8: Example of Mutation for Composite Node

Tree Crossover

During *tree crossover*, the tree is flattened to a list. This enables us to easily *select two random nodes* uniformly from *two different trees*, and swap them.




Figure 9: Example of crossover between two trees.

Fitness Function

As alluded in the theory section, the fitness function f_i is defined as an arbitrary linear combination of game related parameters. The full implementation is described in pseudo-code on the next page.

Algorithm 2 Pseudo-Code for the fitness function

Require: *damage* being the damage given
Require: *armor* being the armour picked-up
Require: *weapon* to be the weapons picked-up
Require: *kills* to be the number of agents killed
 $Score = 5 * damage + 2 * armor + 50 * weapon + 1000 * kills$
if Individual Won **then**
 return $Score + 500$
else
 return $Score$
end if

The score is divided by the number of round each individual has played during a generation. This allows for a given individual to test its fitness against multiple adversaries to get a better estimated fitness, a useful detail when fitness evaluation is done against other peers from the same generation.

4 Results and Screenshots

Below are a couple of in-game screenshots from *Quake 2D*, our testbed for *behaviour tree evolution*. Here we've chosen to show two bots playing against each other in the arena, each of them loaded with a basic hand-made behaviour tree shown in Figure 12.

As can be seen in Figure 10, the environment is composed of *static collidable walls*, *dynamic pick up crates*, *user/agent controlled players*, and *graphical scenery details* based on *tiled textures*. Maps are specified using the *Tiled* editor, which means we can build new environments to test our solution. We've created a serializer for our behaviour trees, that enables us to save/load behaviour trees to disk.

Additionally, we've developed several tools for showing our agent's behaviour in real-time. As can be seen, Figure 11, we're able to *visualize path finding information* for the intelligent agents. Here the *yellow line* is the *target path* specified by the *path finding* system given a *target location* (shown as a *yellow cross*) by the *behaviour tree*. In this example, the target is the *rifle pickup crate*. You'll see a coloured grid being displayed. This tells us something about the state of the *influence map*, where the *blue and red lines* are grids that each agent is able to shoot towards, while *magenta lines* are locations where both agents are able to shoot towards.

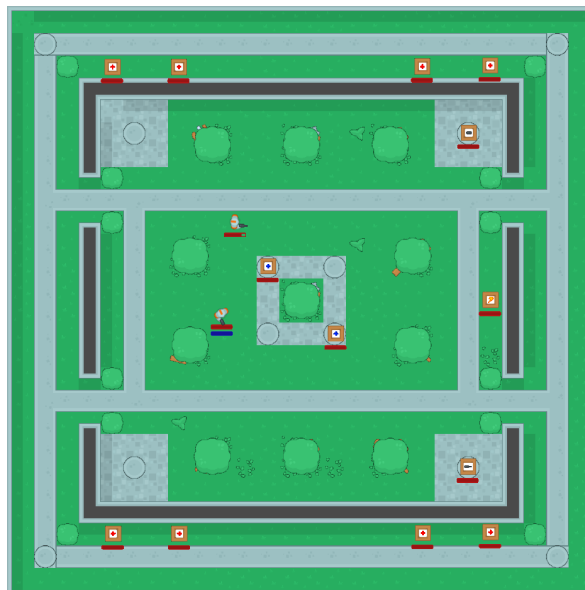


Figure 10: Screenshot of a Bot vs Bot in Quake 2-D

For debugging and analysing agent behaviour, we've implemented a way to visualize BTs, the one in Figure 12 from the Figure 11 situation, in real-time by using a *Graphviz*² library. It opens in a separate window and runs in another thread, where we can see instant changes in the BT's nodes' state.




Figure 11: Screenshot of Path Finding in Quake 2-D

²<https://github.com/nidi3/graphviz-java>

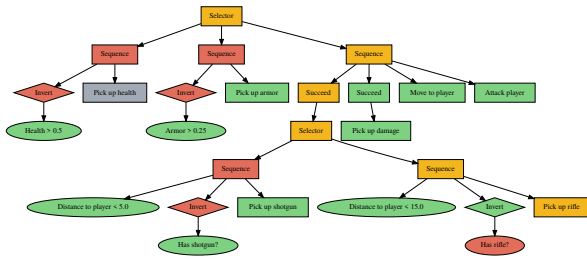


Figure 12: Hand-Crafted BT (Larger in Appendix)

We’ve uploaded a video³ of our solution in action.

Finally, we present the results from our solution. After generating fifty random individuals for \mathcal{P}_0 , we applied genetic programming until we had \mathcal{P}_{50} , the resulting pool of *evolved behaviour trees*. The fitness for each individual in each generation was found by having a match against our crafted behaviour tree, and recording match statistics, used to calculate f_i .

We’ve attempted to evaluate individuals in a generation by pitting them against other individuals in the same generation (we call them *peers* below). Therefore two results will be presented, one where the fitness f_i is evaluated against our *hand-crafted tree* and one against its *peers* in the same generation.

In Section 4.1 we provide two behaviour trees generated with our solution at \mathcal{P}_{50} when playing against hand-crafted behaviour trees. We then give in Section 4.2 graphs of the evolution of the fitness function per generation, for both the peer and the hand-crafted evaluation methods from \mathcal{P}_0 up to \mathcal{P}_{50} . In Section 4.3 we measure the effectiveness of the behaviour tree by looking at the amount of individuals in the generation that have survived against our hand-crafted behaviour trees as the number of generation progresses forward.

4.1 Generated Behaviours

Each fitness evaluation was performed on a randomly selected hand-crafted level to prevent the behaviour trees getting overfitted to a specific level.



Figure 13: Generated tree with a very high fitness score (4750). This behaviour tree is 15 levels deep.

³<https://vimeo.com/242262802>

As can be seen in the image above. This tree is way too big to provide any significant insight, unless interactively evaluated during run time. Therefore we’ve chosen not to provide a scaled-up variant of it. The fitness score is very close to the limit for the levels used during the fitness evaluation.

Here we see an example of a behaviour tree that is quite bloated, which is something we wanted to avoid with our method. But it seems after many generations the trees tend to crossover big sub-trees.

While the depth of this tree is unsatisfactory, the performance it achieves against our hard-coded bot is quite good. It is able to consistently beat it on all different map arenas. In fact, it’s actually hard for even the authors of this paper to beat it consistently.




Figure 14: Generated tree with high fitness (3650). This tree on the other hand is only four levels deep.

Here is a smaller tree of the same generation in Figure 14, it manages to defeat the hard-coded bot on all levels without bloating to unmanageable size. Prior exploration of generated trees indicated that as long as the tree has a *sequence* of *pickup shotgun* and *attack player*, it would reach a fitness of around 1000. Because this kind of sub-tree is quite common, it tends to appear in many BT crossovers.

Since this behaviour tree is a lot smaller, it’s quite easy to analyse the expected behaviour in comparison to the one in Figure 13. At the end of this paper we’ve therefore provided a scaled-up version of this tree so the reader can look through it in more detail.

4.2 Behaviour Fitness

Looking at the general fitness over generations, we can observe that once the behaviour trees routinely starts defeating its adversary, the mean fitness plateaus and the variance fills the entire range of possible fitness values. When evaluating each tree against random peers, we see something similar but it converges with a slower rate. One possible hypothesis of why this happens is that the amount of “junk” behaviours that get by in Figure 16 at the start are all low-performing BTs. Therefore, there is no real competition at the start until one of them gets a good mutation or crossover. After that, this better BT will very likely pass on its genes onto the next generation and generate a lot better BTs than before. In our hard-coded behaviour tree, this happens quickly because the competition at the start is already quite high.




Figure 15: Fitness over generations when evaluating against a hard-coded bot. Higher score is better.




Figure 16: Fitness over generations when evaluating against 2 random peers. Higher score is better.

4.3 Survival Ratio

In the ideal case, the survival rate follows as a consequence of the fitness score. In the two figures below we can see the same trend as in the fitness, the higher the fitness, the higher the survival rate. If this wasn't the case, our fitness evaluation wouldn't provide a good metric for evaluating combat performance. When playing against the hard-coded behaviour tree, it takes around 25 generations until 70% of the individuals in the population are victorious against it. These results then seem to vary between 60% and 90% win-ratio in each generation. On the other hand, when playing against itself, it takes around 28 generations until the same can be said about it. However, it's actually able to achieve a high maximum win-ratio, and has less variation between each generation, usually staying between 76% to 95% win-ratio.




Figure 17: Victories for each generation when evaluating against a hard-coded bot. Higher value is better.




Figure 18: Victories for each generations when evaluating against 2 random peers. Higher value is better.

5 Discussion and Outlook

There are still some areas in the project which merit further analysis. We'll start with the results section, and see what we can extrapolate from the fitness and survival scores. Afterwards, we look at some limitations of our solution, compare it to related work, and then end with some future improvements.

Convergence Rate

As can be seen in the plots in the result section, the evolutionary algorithm converges to a high game proficiency in just a couple of generations, and manages to stay that throughout the evolution process.

The fact that it manages to do this may reflect poorly on our skills as game designers. Since one expects that the more complex behaviours that are required to proficiently play the game, the more generations it will take until it produces such well-tuned behaviour trees. As we mentioned before, all it takes to beat our hand-crafted behaviour trees is to play very aggressively. And the evolved trees which are generated clearly reflect this fact.

It can be viewed as the fitness function, tree mutation, and pool of possible actions being well tuned to the specific requirements of this project. Which interpretation is correct is left to the reader.

A possible explanation to the different convergence rates when evaluating against a hard-coded tree or adversarial peers is the stochastic nature of both level selection and the randomized adversary. The adversarial peer should be much closer in performance compared to the hard-coded tree. Once the trees which can defeat the hard-coded bot have been generated, additional fitness can only be reached by overfitting to the given level. The trees with peer evaluation on the other hand, can't find a solution on how to reliably defeat the enemy. The adversary co-evolves with them which makes it very hard to find a "holy-bullet" behaviour since almost all individuals in the next generations will have it.

When tested against a human player the evolved behaviours becomes very hard to defeat. Even for behaviour trees with a fitness score of around 3000. One of the reasons could have to do with the agent's superhuman accuracy when shooting, but it has to do with how greedy the agent is when it comes to pickups, sometimes it can play really safe, and keep refilling it's health and sometimes denying pickups!

Limitations

Since we decided to have high level actions, the tree's behaviour often appears clunky compared to a human player. As an example, two bots might decide to walk to the same health pickup when both have a very low health instead of just shooting at each other. Current fitness evaluation is bound to the render thread, which severely limits the number of evaluations that can be done during one second., the amount of actions, object types, and levels have been fairly limited in this project to produce faster results and ease debugging. This leads to the BTs having a limited repertoire of behaviours.

Related Work

As we mentioned before, our genetic programming solution was initially based on the work by *Colledanchise et al.* [3]. The solution proposed in the paper did not pan out in our case since they use an selection heuristic during the mutation phase, allowing for higher fitness increase over fewer generations. The paper used primitive actions which would not scale to the type of game used in this project without *very* large behaviour trees. With this said, it seems like low-level actions might produce more "clutch" behaviours in this kind of fast-paced games.

Also, even though the behaviour tree in Figure 13 is not unusable in practice, it still became deep. The bloat control and pruning steps in *Ögren et al.* [3] might not have been such a bad idea. However, our solution is algorithmically simple, and gives acceptable results. Thus, we recommend starting simple, and then follow *Ögren et al's.* [3] steps when needed.

Future Work

Some improvements are to extract the behaviour tree and genetic programming parts as an independent library and to use heuristics during the creation of each new generation. To add additional behaviours and mechanics to the project and have it compete with itself instead of a hard-coded adversary, which has been tested to some extent during the project but not enough to yield interesting results. It would be good to run fitness evaluation in headless mode and become physics independent.

If you are interested in our solution, feel free to look at our code and apply some of the additions. ⁴

⁴Repository: <https://github.com/sci10n/Quake2D>

References

- [1] N. A. Barricelli et al. Esempi numerici di processi di evoluzione. *Methodos*, 6(21):45–68, 1954.
- [2] A. J. Champandard. The behaviour tree starter kit. *Game AI Pro*, pages 73–91, 2014.
- [3] M. Colledanchise, R. Parasuraman, and P. Ögren. Learning behavior trees for autonomous agents. *arXiv 1504.05811*, 2015.
- [4] M. Dawe, S. Garolinski, L. Dicken, and T. Humphreys. Behavior selection algorithms: an overview. *Game AI Pro*, pages 47–60, 2014.
- [5] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.
- [6] J. W. Hoff and H. J. Christensen. Evolving behaviour trees: Automatic generation of ai opponents for real-time strategy games. Master’s thesis, NTNU Dept. of Info. Science, 2016.
- [7] D. Isla. Managing complexity in the halo 2 ai system. In *Proceedings of the Game Developers Conference*, volume 63, 2005.
- [8] C.-U. Lim, R. Baumgarten, and S. Colton. Evolving behaviour trees for the commercial game defcon. In *European Conference on the Applications of Evolutionary Computation*, pages 100–110. Springer, 2010.
- [9] D. Mark. Modular tactical influence maps. *Game AI Pro*, 2:343, 2015. [Online; 2017-11].
- [10] R. Nystrom. *Game programming patterns*. Genever Benning, 2014.
- [11] J. Orkin. Three states and a plan: the ai of fear. In *Game Developers Conference*, volume 2006, page 4, 2006.
- [12] D. Perez, M. Nicolau, M. O’Neill, and A. Brabazon. Evolving behaviour trees for the mario ai competition using grammatical evolution. *Applications of evolutionary computation*, pages 123–132, 2011.
- [13] Simpson, Chris. Behavior trees for a.i. , 2014.

