Racing Artificial Intelligence

Tim Beaudet – December 2nd 2009

The Problem

Most racing games these days have several problems with their Artificial Intelligence, AI, behaviors. A large category of these problems come in the form of the AI cheating. There are several ways for the AI to cheat; catchup, different physics, super control, and super knowledge. Some of these are perfectly fine for use in game AI. But what about a simulation, where these super natural AI algorithms give a disadvantage to the player who needs to react on the estimation of the track reference points, tire traction and other obstacles.

More things to look into is how dependent the AI actually is on the physics system. Then we have basic car control, keeping it at the limits and regaining control after breaking the limits. Human drivers make many mistakes every lap, but they are able to keep control of the vehicle and stay on course. And finally driver behavioral patterns; when to overtake, which direction to overtake will be described. All of these points will be looked into further as we search for the problem with taking AI from a game and applying it for a simulation.

First issue is dealing with cheating techniques like catchup, different physics, super control and super knowledge. Each technique will be examined to show how it is used and how the AI driver benefits from it, or how the player falls to a disadvantage. Finally a possible solution will be thought up, so that both the player and the AI driver stand on the same level field. This is not a section where detail is spent on fixing the problem, only general ideas of what to do instead. If the general idea has risk, or parts that could go wrong it will also be discussed in the section. So the first issue on the board is the famous catchup algorithm.

Catchup

Catchup is where the AI gets a boost in speed when the player is ahead by a certain percentage. Alternatively when the player is behind, the AI drivers slow down, make mistakes or the player gets a boost in speed to catch back up to the AI.. It looks like both AI and player are evenly matched at first thought. After all, if the player falls behind they get a boost in speed, or some other mechanism to catchup, just like the AI would. While this technique is good for an arcade racer, you do not see real racers slowing down and making mistakes because a driver fell behind them. To put this AI helper into a simulation would break realism, as a car, AI controlled or player controlled, can only go so fast, and adding a speed boost of any kind will break the rules of a simulation.

So even if the catchup algorithm keeps the AI driver evenly matched with a human driver, it goes about doing it by breaking rules of a simulation. This will perhaps be an easy fix, just don't include the technique in a simulation, but that presents the issue of multiple skill levels. Games that use catchup have difficulty customized to the players ability. In the true simulation, it doesn't matter that the player is over challenged or under-challenged, however in games the challenge for the player is extremely important. The problem with removing this from our simulation is the loss of challenging the player at their own skill level. That is a reasonable sacrifice for now.

Different Physics

The next AI advantage technique is when AI drivers run on a different set of physics than the player car. This is a little harder to detect as it can feel similar to the super control algorithm described next. While entering a corner the player accidentally slides into an AI driver, the player loses control while the AI continues like nothing happened. Some methods of this technique give the AI more mass to keep their car more stable under such actions. Sometimes they use the collision algorithm to give the AI car more advantage, and sometimes it is not as noticeable as that example.

The problem here could be that the AI car has more power, less mass or a different controller scheme. Where the players control needs to enter their steering wheel, and add forces at the tires, the AI controller may just turn the car by rotation, eliminating the need of car control or bypassing the physics of the tires. This form of AI will likely tie the AI and physics system closely together. After closely examining the situation you will find the player at a disadvantage, although it is less noticeable than the previous example with the catchup techniques.

The player relies on the physics system to keep the car under control. If it was as easy as rotating the car without worry about about front end losing grip the player would have no worry about understeer or oversteer.

The solution here is quite simple. Taking physics completely away from the AI driver, removing the ability for the AI driver to know the exact information about a collision that happened to their car. Giving the AI driver the same controller options as the player; throttle position, brake position, steering and shifter. This brings up another part of the overall problem, which will be discussed in more detail later, the knowledge of the physics system.

Super Control

Sometimes you are side by side with an artificial driver, and after slight contact you start to lose control and slide around, while the other driver magically continues. There have been observed cases where the AI does not use a different control of physics, but their input reaction is so quick and precise that their car regains grip before you even noticed the bump. This lightning fast, smooth response is unlikely in the human world, and is what we are calling the super control phenomenon. While some skeptics can argue the driver saw it coming, this probable happens less than ten percent of the time. Also in the event of preparation, artificial driver should be focusing on keeping control over the car more than focusing on entering the next corner at a perfect point of entry.

So the advantage here is quite easy. Go to a steering wheel, and spin it 180 degrees, it takes a bit of time. Not minutes, but still time that matters. If you could turn the steering wheel like the AI can with a super controller, then it would happen instantaneously. Another good point here is in shifting. With a standard “H-shifter” or even a bump “shift lever” transmission the driver needs to move their hand from the steering wheel to the shifter, takes at least one tenth of a second or quite a bit longer to do. While a hand is on the shifter, you can't turn the wheel 180 degrees either without the delay moving the hand back. These are things AI with super control don't worry about.

To fix this issue the AI will be restricted to mimicking a humans reactions as closely as possible. Delays between shifting will be necessary unless the driver prepared to shift, which means they placed their hand on the shifter. Any racing AI algorithm that has super control, but no catchup an no special physics could benefit from this solution. The driver would be able to choose the exact throttle position, steering position or otherwise, and a layer over that would slow the inputs as required.

Super Knowledge

The final advantage the AI driver gets over the player is by collecting much more accurate information about the track. As a human we get a certain field of view that we can see. Yes, that means your mom doesn't have eyes in the back of her head like she claims. So the AI driver should deal with this restriction as well. A human is good at judging their speed and distances to objects closer to the center of the field of view. These judgments are less accurate when it comes to peripheral vision, off to the edges of the field of view. Also another limitation to the player is objects hiding reference points. A player can use a cone to indicate a brake point, or turn-in spot for a corner, when a car is in front of them the player can't see the cone, and will need to judge their brake spot from less accurate points of reference like a billboard further away.

This disadvantage is easily noticed, as the AI knows where everything is at all times. Even if a hill is blocking the view. The solution starts getting more complex here, based on some existing algorithms the AI knows the left and right boundaries of the track, and the direction they should be traveling, as well as where the best line is located in that area. Sometimes these points tell the driver acceleration information and other things of interest. This can be important information to how the driver follows the track. The solution proposed give the driver the same reference points that the player would get. Actually due to computation time, editing time, and other various things it might be less reference points, which is one limitation and possible point of failure.

So now the driver knows with some level of estimation where the reference points are in relation to the driver. In the more technical, implementation idea section the problems and solutions will be discussed in greater detail. It will likely need to get down and dirty, involving math and parts of code. In order to keep this section general, we can say that giving the driver only the information a player can get will be challenging to do in technical terms. It is also to mention that this could cause the driver to be less challenging, and have a hard time to go around an unknown track.

If this point fails, it may be possible to use the existing, tested and proven method for getting the AI driver to circle the track, but then combine the two systems for an added benefit. This would help with the realism of the driver because he would not be able to see the cars behind him with a very accurate estimation unless he was checking the mirrors, which would blur the reference points in front. Most of the effort will be placed in getting the driver to know what to do from these reference points, attaching instructions or memories to these points and having the driver use them to drive properly.

So to summarize the above one major problem with AI in simulations is that the AI has advantages over the player. Removing these advantages can be easy in some cases, and rather difficult in others. Removing catchup and different physics for the AI driver will be easy while developing the system, but doing so can lead to AI that is not so challenging for some drivers, which is an acceptable side effect since this is a simulation and not a game. The main area of focus is removing the superior knowledge about everything on track and the superior controls that AI currently have in most algorithms.

Now lets shift gears from the cheating techniques to other possible problem areas, like the dependency of physics in AI. As a player you don't know the physics of the world you play in. Some things are assumed, you will lose traction and slide at times. The AI shouldn't need to know anything except that tires have lost grip. Players can get this information based on sounds, and to some degree physical changes. In forces from the steering wheel and a sense of rotation among the balance mechanism in our brain. With the driver detached from the physics simulation it will require only bits of information that will be sensed.

It is impossible to make a driver race a track without them knowing how close to the limits their tire is. Also important that the driver feels the acceleration and rotation forces, but these should be collected by a sensor rather than being directly dependent on the physics. In theory then, the AI driver should be driving to the limits of the physics system provided, regardless of what changes happen. Which would allow the system to be reused. Basic car control does, to a very small degree hold a certain knowledge of the rules of physics.

Basic car control doesn't depend on the physics, but does make assumptions and will create a few problems if these assumptions were not true. Basic car control is about keeping the car under control or regaining control. Basic ideas and rules are assumed: When understeering, figure out why. Are you pressing the brakes too hard? Release them a bit. Turning too much? Turn less. Going too fast? Apply brakes. These certain rules are taken for granted by humans, and some games may have oddball physics where these rules make the situation worse, but this is an assumption the AI needs to make. This area does not add anything to finding the problem, but brings up the limitation of the previous assumption.

Human drivers make mistakes, lots of mistakes, many times per lap. Actually from reading several books, and playing a few racing simulators it is quite accurate to say a fast racer knows how to detect a mistake and fix it faster than a novice racer. The hopes of this AI simulator is not to make a perfect AI driver, but to make the driver make, detect and fix mistakes based on judging their distances to reference points and the other information available. This is probably the most critical part of the AI simulation. Especially when passing and dealing with other cars is concerned.

So by this point if everything is correct, the driver should be able to driver laps around a track. But that would get quite boring, we need to add additional cars and race! But this posses another wide set of problem. Current AI systems likely have each driver know about each other driver, and not just know where they are but also what they are thinking and planning to do. Humans don't get this information from their competitors, they just use their observations to see that a move might be made.

These observations will need to be filtered by the drivers field of view, just like the other reference points. A driver should be able to check their mirrors for more accurate positioning of a driver behind them or to the side, but without this check it should remain an estimation. Deciding when and how to pass could bring up issues. From looking into other AI algorithms it seems the use of an ideal line and a passing line is typically used to tell the AI when, where and how to pass. This is not a terrible idea as it simulates the general knowledge a player has.

A player knows certain areas are bad for passing, and they also know the basic ideal line. Both of these lines would not contain information that tells the AI how to change any input. While this is not a terrible idea, it is not a great idea either. Currently it is substituting for the lack of a better idea. The problem with having these two lines is the AI drivers will typically pass in the same areas. Passing and dealing with other cars is certainly a huge risk currently, and the fallback path is using the defined driving lines as hints of when, where and how to pass.

So after all of that what is the actual problem? Giving the AI control of the physics in any manner is certainly one issue. After removing the physics dependency from the AI there will be no form of catchup or different physics control for the AI. The AI will use senses just like the player to plan what to do. The AI having instantaneous reaction times with the controls, and knowledge of everything on the track, even when it is out of site is another issue. That issue leads to the AI becoming too perfect, using exact information to know what to do with the controls and following through with it.

If you have anything you would like to add, see a mistake, or think this project should take into account please e-mail me at:

Tim Beaudet
blackbird04217 at yahoo dot com