EDUCATED PLAY!

Galactic Arms Race

EVOLVING THE SHOOTER

THE STUDENT-BUILT GAME

GALACTIC ARMS RACE is the intriguing result of ongoing work from the Evolutionary Complexity Research Group (EPlex) at the University of Central Florida. Starting with the basic idea of an online multiplayer space shooter, GALACTIC ARMS RACE adds a unique wrinkle to the competitive formula by featuring particle-based weapons that evolve into novel configurations during play.

We contacted Kenneth Stanley, the team's faculty advisor, to find out more about the automatic content generation driving GALACTIC ARMS RACE.

Jeffrey Fleming: *Can you tell me a bit about the team that worked on GALACTIC ARMS RACE (GAR)?*

Kenneth Stanley: The team reflects the origin of the game inside a research group at the University of Central Florida. I supervised the project as its faculty supervisor and my Ph.D. student Erin Hastings took the lead in software development and technology integration. The project required integrating novel AI technology developed for the project into the game. The rest of the team was rounded out by volunteers who were mostly undergraduate students interested in gaining experience working on a game. Overall, the project represents a major volunteer and educational effort driven by people's passions, with little financial support.

JF: *What tools did the team use to create GAR?*

KS: GAR is made in XNA. It also uses NEAT and something called "NEAT Particles," which is a technology developed before GAR to allow NEAT to evolve particle systems.

JF: *What is the idea behind the NEAT algorithm?*

KS: NEAT stands for NeuroEvolution of Augmenting Topologies. I invented NEAT at the University of Texas at Austin when I was a Ph.D. student working with my advisor Risto Miikkulainen. As its name implies, it evolves artificial neural networks, which are kind of like little artificial brains. The innovative aspect of NEAT is that the brains it evolves actually get bigger as evolution progresses, which is what the word "augmenting" means in its name. In simple terms, the implication is that behaviors can become smarter and more complex over time.

NEAT is the core of the algorithm that evolves the weapons in GAR. Actually, for GAR we introduced a variant of NEAT called cgNEAT, which stands for "content-generating NEAT." Believe it or not, a neural network evolved by cgNEAT drives every particle in every weapon in GAR. So the neural networks are actually controlling the way weapons behave. Because the weapons are evolving through NEAT, their behavior can become more complex and intricate over time.

JF: *How does cgNEAT decide which weapon to evolve in GAR and which are dead-ends?*

KS: The way cgNEAT works in GAR is that it tracks which weapons people like by observing which ones are fired the most. Those that are popular become the "parents" of new weapons that are spawned in the galaxy. Thus the question of which weapons evolve is answered by which weapons people like. If people like them, cgNEAT makes new variations of them and spawns them in the galaxy.

JF: *Are the evolved weapons specific to a single instance of the online game or are they part of a persistent world?*

KS: In multiplayer mode, the evolved weapons are stored on the server, so they generally persist as long as the server. In that sense, they are part of a persistent world for each server. So the interesting situation is created in which evolution can continue over months or years.

JF: *Is the neural network very processor intensive?*

KS: In GAR, every single particle from every evolved weapon is controlled by a neural network and even when there are ten people on the screen at the same time all firing different weapons at once over a network, GAR players will not experience any slowdown. So from that perspective, today's CPUs are more than capable of handling many simultaneous neural networks being activated at the same time. Neural networks tend to be compact and require only a few floating-point operations, so they can often be less computationally expensive than more traditional control schemes. However, of course, if neural networks are allowed to grow very large, they can start to be more expensive. Yet that size is well beyond what is needed for the type of control in GAR or many other games.

JF: *How difficult was it to integrate online play into the game?*

KS: Integrating online play was a challenge because we had to get the system to perform evolution over the Internet, which means that genomes and fitness information literally have to be sent back and forth through messages over the network. There is not much precedent for a real-time Internet-based evolutionary system like that. For example, if a player flies into your view with a weapon you've never seen before, the neural network for that weapon must be transmitted to your computer right away so that the other player's weapon looks the same to you as it does to the other player. However, once the proper information is set up to transmit to the right places, the overall evolutionary algorithm works seamlessly and is not hard to manage.

resources

GALACTIC ARMS RACE
http://gar.eecs.ucf.edu

Evolutionary Complexity Research Group at UCF
http://eplex.cs.ucf.edu/publications/2009/hastings.ieeetciaig09.html

Google Groups: Procedural Content Generation
http://groups.google.com/group/proceduralcontent

Search-Based Procedural Content Generation
http://julian.togelius.com/Togelius2010Searchbased.pdf