

## Before you walk...

It's generally accepted that real humans usually get around by walking on two legs. This is widely held to be the most efficient form of motion for a bipedal creature, or at least the most desirable. Maybe we've evolved to keep our hands somewhat cleaner than our feet, or our brain is capable of better balance than our quadruped cousins. Perhaps it's simply that everything except our legs is really very weak. Whatever the reason, you don't often see adult humans moving around by hopping on their knees, doing handstands, or launching themselves into the air by smacking their head against the ground.

Walking, though, is a somewhat complicated task. It takes humans more than a few months of constant study to learn how to toddle, and people constantly adapt their walk throughout the rest of their life to cope with changing bone length, changing muscular strengths, and injuries.

Studies of how humans walk are conducted in many disciplines: animations, robotics, bio-engineering, dance, martial arts are just a few examples. Teaching a computer how to walk is a complicated task, and sometimes you have to learn to crawl before you walk!

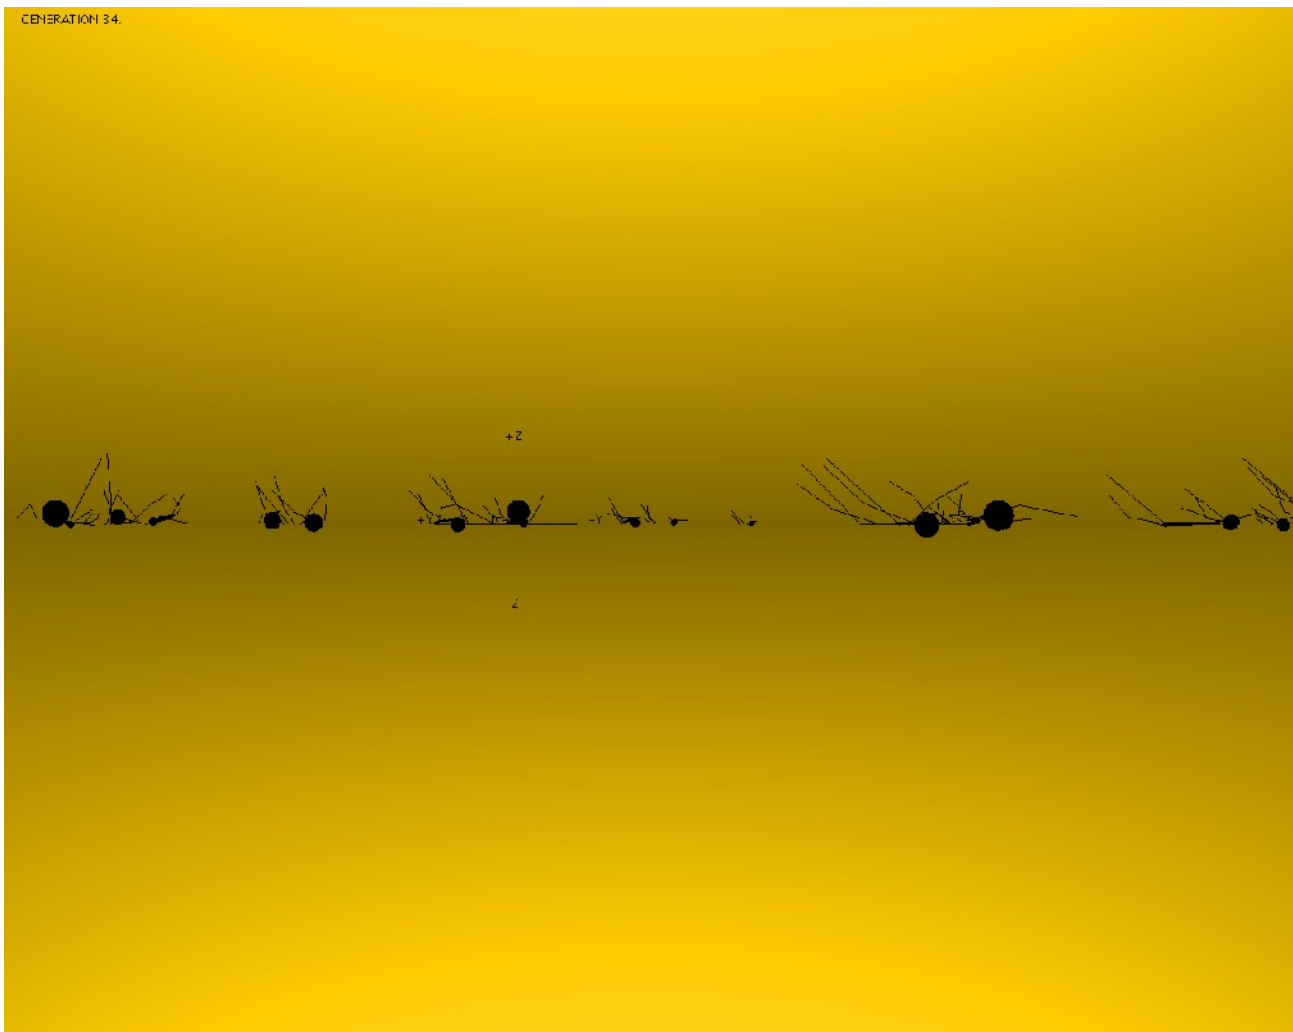


Figure 1: Stick figures crawling.

## ...you must crawl!

The goal of this project is to attempt to create human-like crawling motion in a simplified physical environment through standard AI methods. In particular, the code included with this paper (which generated the figures pictured here) uses a genetic algorithm to allow successive generations to "learn" what kind of limb motions are important to walking. The individual stick figure people (endearingly called "stickles") are simple human-like beings that cannot see, hear, or feel. They are generally pretty flexible, and extremely strong. Finally, they have no concept of pain or fatigue, which naturally influences the way they learn to crawl.

## The Stickles

Stickles have only ten functional bones: torso, left upper leg, left lower leg, right upper leg, right lower leg, left upper arm, left lower arm, right upper arm, left lower arm, and head. Between each of these bones is a joint, which is restricted in motion based on the joint. These restrictions represent the limits of average human ability: for example, most humans cannot raise their leg straight backward perpendicular to their spinal cord. Stickles cannot do this either. However, joint limitations on stickles are completely independent of the positions of other joints. As an example, a stickle can raise its knee to its chin, as most humans can (or close to it). It can then also extend its knee so that its leg is perfectly straight, while maintaining the fact that its knee is touching its chin. Many humans (or at least the author) cannot do this!

Stickles have no model of fatigue or limited strength. Each joint is controlled independently across its legal range by a function  $F(t)$ , where  $t$  is the current elapsed time in the simulation.  $F(t)$  is the sum of a number  $N$  of sine functions, each of which has a different frequency. These component functions may also have different shifts and amplitudes. As a result, all movement is periodic given a long enough simulation.

## The Learning

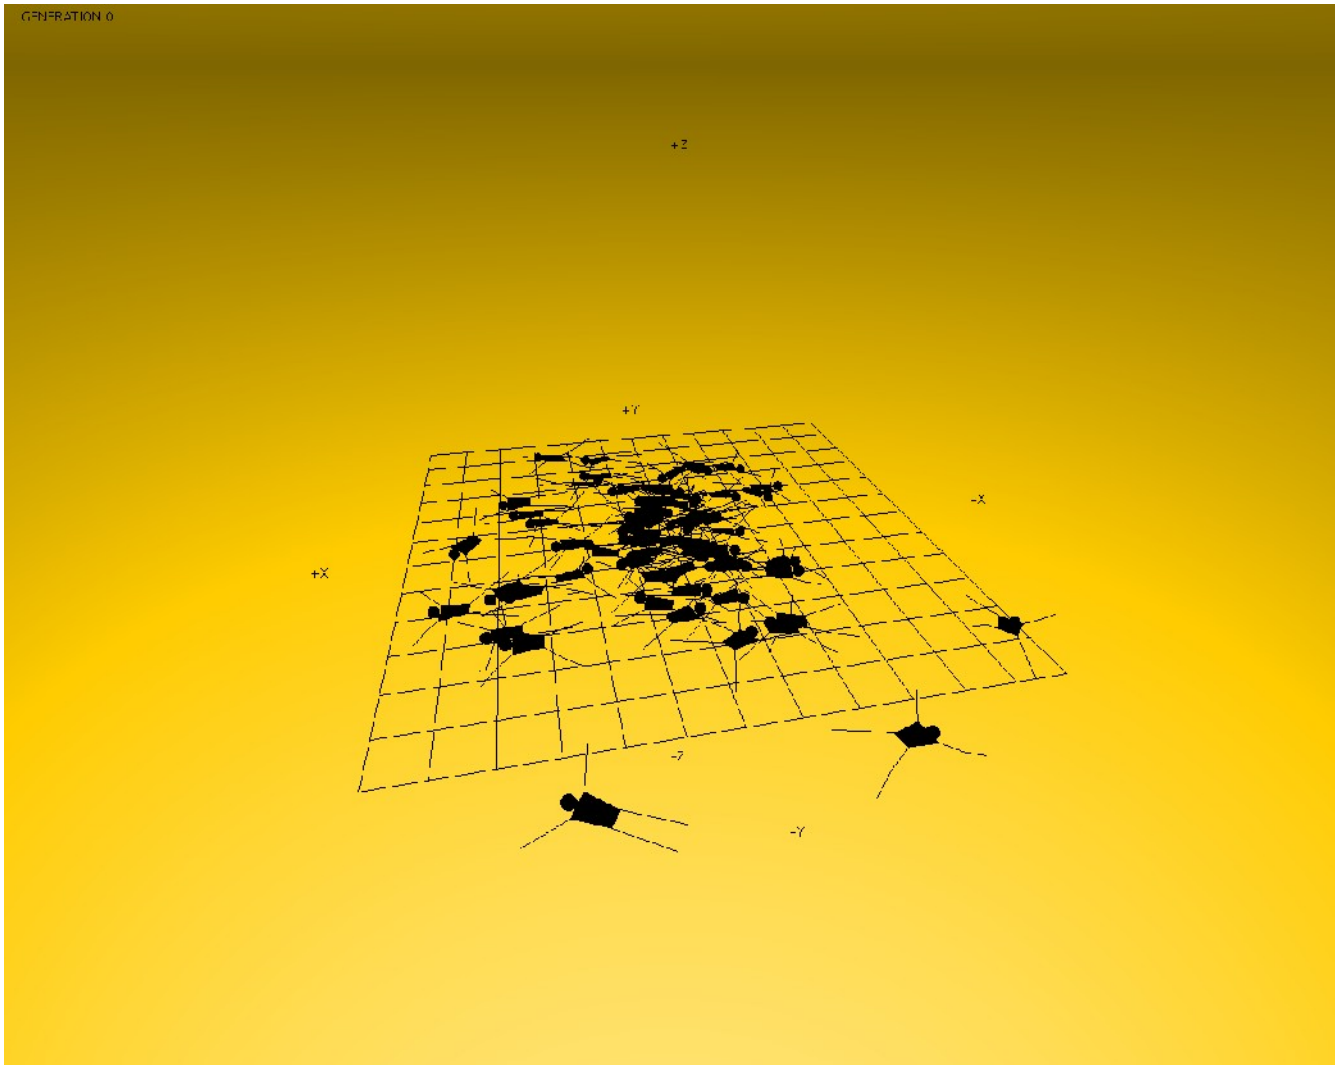
Stickles operate in an external environment called Stickleworld. Their motions act upon the environment, and the physical rules of the environment determine how far they travel. At the end of every simulation period (15 seconds at 60 frames/second, longer at slower framerates - updating is done on a fixed clock, not the system clock), stickles are awarded scores by how far they move in the positive X direction. The top four scorers are selected, and produce 32 children by the following breeding algorithm:

Every possible pair (including stickle1 with stickle1, stickle2 with stickle2, etc) produces 4 children. Two are "left-bred", the other two are "leg-bred". Left-bred children inherit the motion of their left arm, left leg, and head from one parent, and the motion of their right arm and right leg from the other. Leg-bred children inherit both arm and head motion from one parent, and leg motion from the other parent. There are two of each kind of bred child, since there are two ways to arrange each pairing. For example, one leg-bred child inherits parent1's arms and head, and parent2's legs. The other leg-bred child inherits parent2's arms and head, and parent1's legs.

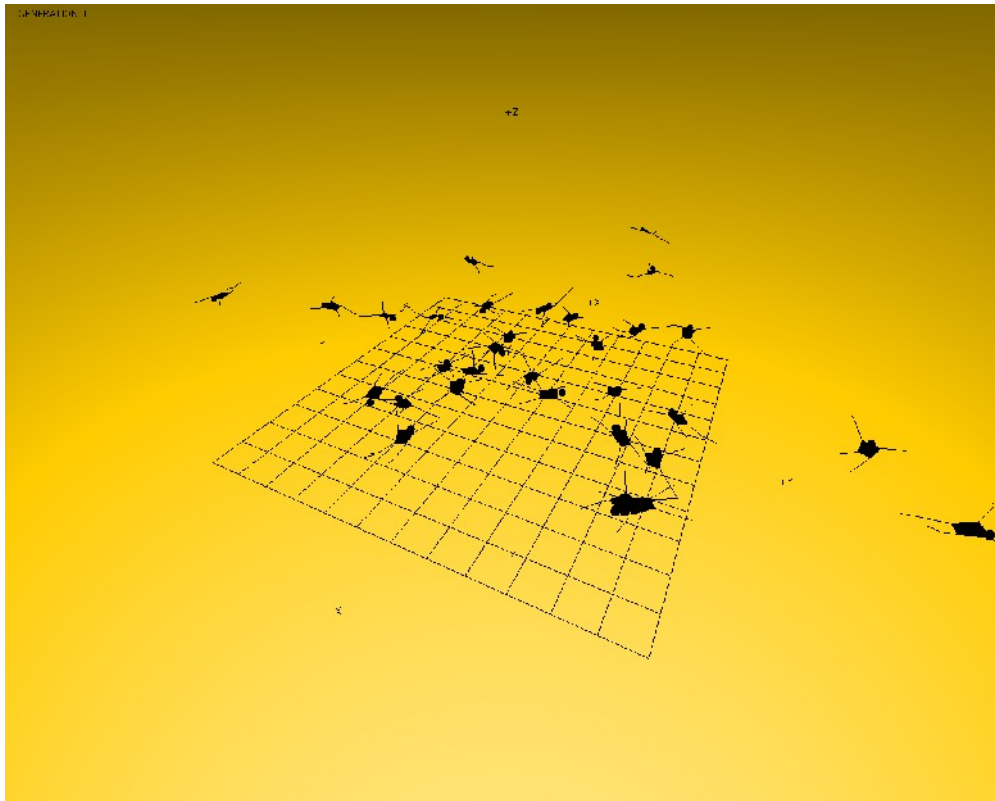
All children are then mutated, allowing the shifts and amplitudes of all component sine functions

of all joints to change by up to a fixed mutation scalar (provided as a command line argument). Successive generations of children are bred.

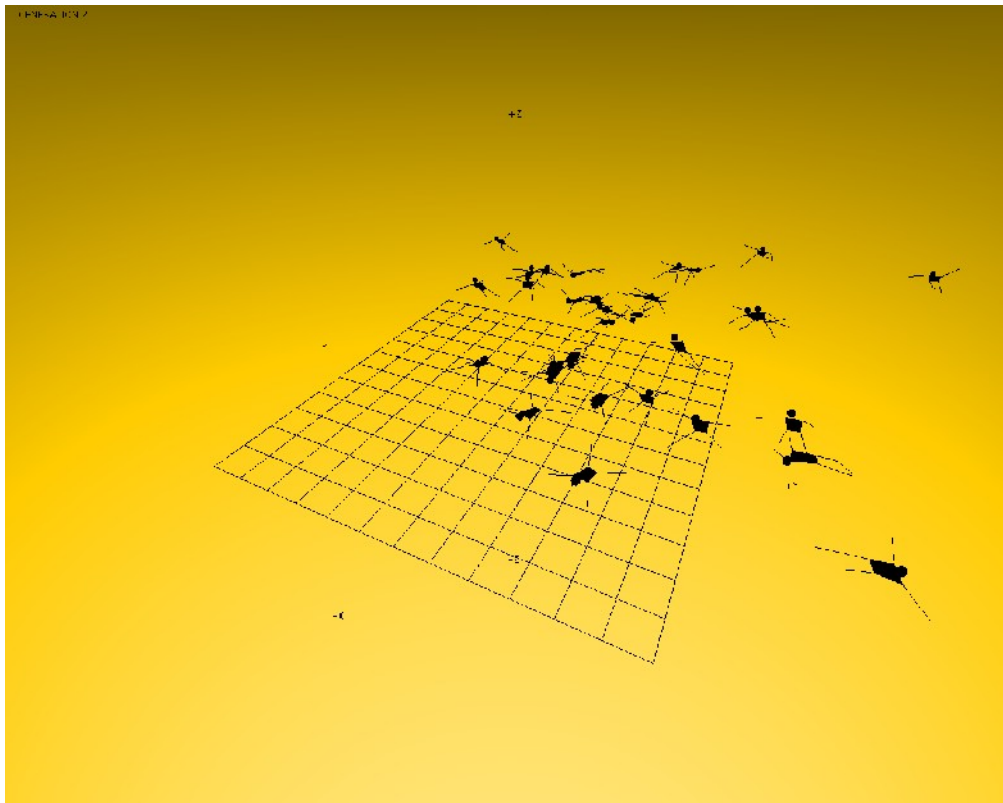
Stickles do not intersect each other, and so they may pass through each other freely. They all begin at (0,0), standing upright. After immediately falling over, their motions usually result in some kind of crawl. Fast crawling in the proper direction is encouraged, and so eventually they improve upon their previous scores.



*Figure 2: Generation "0" of one simulation: all stickles move limbs randomly. All stickles have started from the centre of the grid.*



*Figure 3: End of generation 1 (same simulation as prev. figure), note more stickles have moved further, and the bias of motion toward the +X direction.*



*Figure 4: Generation 2. Note further motion in +X direction.*

## The Physics

The physics of the simulation are somewhat irrelevant. The stickles are completely unaware of how their motion fits into the rules - they simply know what the physical world rewards with motion and what it doesn't. They will, therefore, adapt to changing physics.

However, the physics of the world used in this study does influence the results of this study in several important ways, and so it is important to understand.

The physics of the Stickleworld follow two basic rules: no stickle may jump (have no points of contact with the ground), and no stickle may fall through the floor (have one or more vertexes below  $Z=0$ ). As a consequence, this means that a theoretical stickle sitting with legs straight ahead could "snap" them down while leaning forward and would stand straight up: while the stickle would fall backward slowly (in real life, a human would slam his back into the floor), it can use the fact that it is not allowed to fly to bend some of the "real" rules. This is most obvious in the "highkick-knee-crawl" that evolved midway through the project, before stickles were allowed to rotate (and thus fall down).

In the highkick knee-crawl, a stickle kicks a leg as far forward and up as possible, while stretching the back leg as far back as possible. As the front leg begins to come down, the knee at the back leg is bent, raising the back foot. Since stickles must have at least one point touching the ground, the stickle falls to its back knee until the front foot comes down, which becomes the new point of contact. Since it continues to descend, the stickle rises off the back knee, and brings the leg forward. As the knee is still bent at this point (but beginning to unbend), the leg passing forward does not touch the ground, and so does not disrupt the stickle's forward motion. The stickle is propelled forward by the other leg, which snaps backward and becomes the hind leg. The process then repeats.

The ground is treated as having an infinite coefficient of friction when the stickle has only one point of contact. There is a small layer of "cushion" tolerance that the stickle can interact with: any vertexes beneath  $Z=T$  for a small but positive  $T$  are considered "touching" the ground for the purposes of friction. Thus, if the stickle has only one point of contact with the ground, regardless of how the stickle moves, that point will be fixed: the stickle may propel himself by moving that point *relative to his own center of mass*.

Things become more complicated when a stickle has multiple points of contact with the ground: in this case, slip at each point is minimized by taking net motion to be an average of the relative motions to the points of contact.

A stickle with one or two points of contact with the ground is almost always falling. If there is only one point, the entire body of the stickle is rotated around an axis perpendicular to both straight up and the vector from the centre of mass to the point of contact. If there

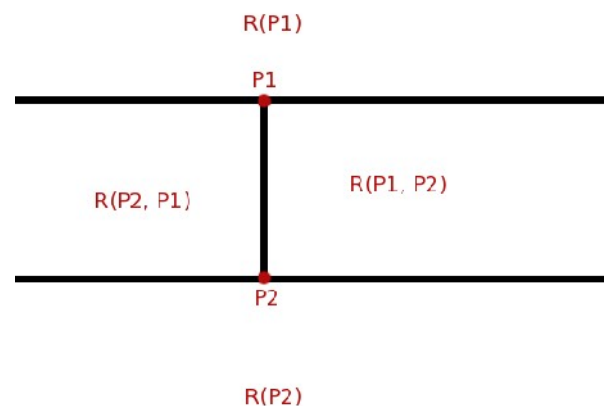


Figure 5:  $P1$  and  $P2$  are points of contact. If the projection of the centre of mass onto the horizontal plane lies within region  $R(P2)$ , the stickle falls as if  $P2$  was the only point of contact. In  $R(P1)$ , the stickle falls as if  $P1$  was the only point of contact. In the other two regions, the stickle falls toward the centre of mass around the axis defined by the line segment between  $P1$  and  $P2$ .

are two points of contact, the stickle may either topple sideways over one point of contact, or fall on the axis defined by the two points of contact, as explained in Figure 5.



*Figure 6: Stickle performing the "highkick-knee-crawl", at the peak of the highkick moment. Note that this is a different physical world than all "yellow" simulations (hence the red background): here stickles' torsos are not permitted to rotate in any direction. This behavior evolved in almost every simulation under these physical rules. Arm motion is often irrelevant, since the arm behaviour never evolved to touch the ground (though they could reach if both legs were near horizontal).*

## The Results

In all simulations run, the learning was very apparent, though many times the stickles did not exhibit human-like motion. A qualitative description of several interesting motions follows:

All conducted with mutation rate .15, initial diversity .15, 64 initial beings. NS means "number of sines per joint" - the number of component sine functions.

NS= 3: Legs evolved to become uninvolved - purely an arm crawl, in a cyclic swimming motion.

NS= 3: Sweep left leg out and then under body - kick off it to get a fast motion that tips back side up, then wheel-barrow walk on both hands for as long as possible.

NS= 3: Sweep right leg out then under body, kick off it, and then drag self forward with hands until leg can come around again.

NS= 2: Sort of an inverted crawl, face up and legs first.

NS= 2: Crawling forward: right leg evolved to stay raised and out of the way, left leg cuts under and pushes forward, arms constantly swim forward

NS= 4: Right leg held aloft: useless. Arms windmill in crawling motion, synced so that left leg swings in (knee unbent) under body, then pushes off.

NS= 4: Crab-walk involving flopping onto chest, pushing oneself up with arms, scrabbling sideways, falling onto chest, etc.

NS=15: General contortions, two flails forward one failing motion back back. Movement generally escapes description: it's frenetic and complicated.

NS=15: As above. Whatever that motion is, humans shouldn't do it. Involves a lot of leg intersection.

Best Movement  
(by Generation)

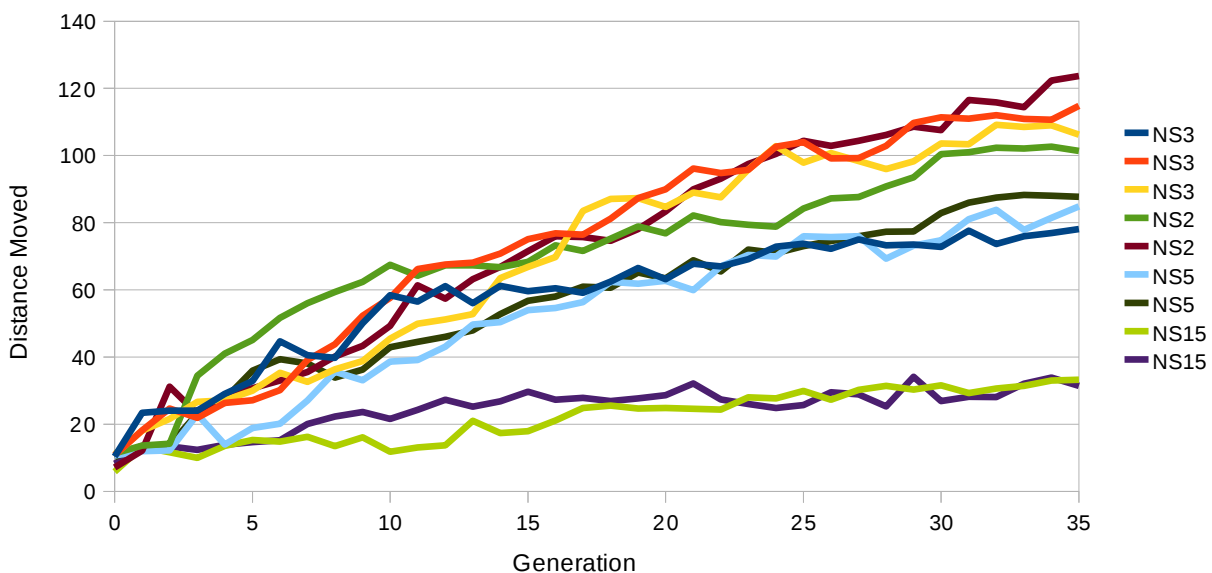


Figure 7: Graph of maximum performance of above qualitative motions.

Even in the extremely complicated motion cases (NS15s), it's clear that there is some improvement over time - however, the relative impact of mutations may be too great, and all motions too complicated to really decide what's helpful and what's not. The simpler, more sinusoidal movements associated with only 2 or 3 component sine functions perform (and learn) much faster. This intuitively makes sense: a brief trip around the room and some introspection reveals that hip and knee motion is not jittery.

In any case, the fact that the algorithm learns a worthwhile motion given the physical system is obvious. A few other correlations were found: lengthening limbs creates generally faster motion, while chopping limbs off results in slower motion. However, asymmetric stickles (with one long leg or arm) do not perform noticeably better or worse than symmetric stickles (assuming the symmetric version has limbs as long as the longer limb of the asymmetric stickle).

Included below are some sample images of evolved motions, though still images are relatively worthless in communicating how the motion works: it's probably a better use of your time to simply run the program once or twice to see how the different motions evolve. Alterations of the stickle.anim file in the data folder will change the geometry of stickles.

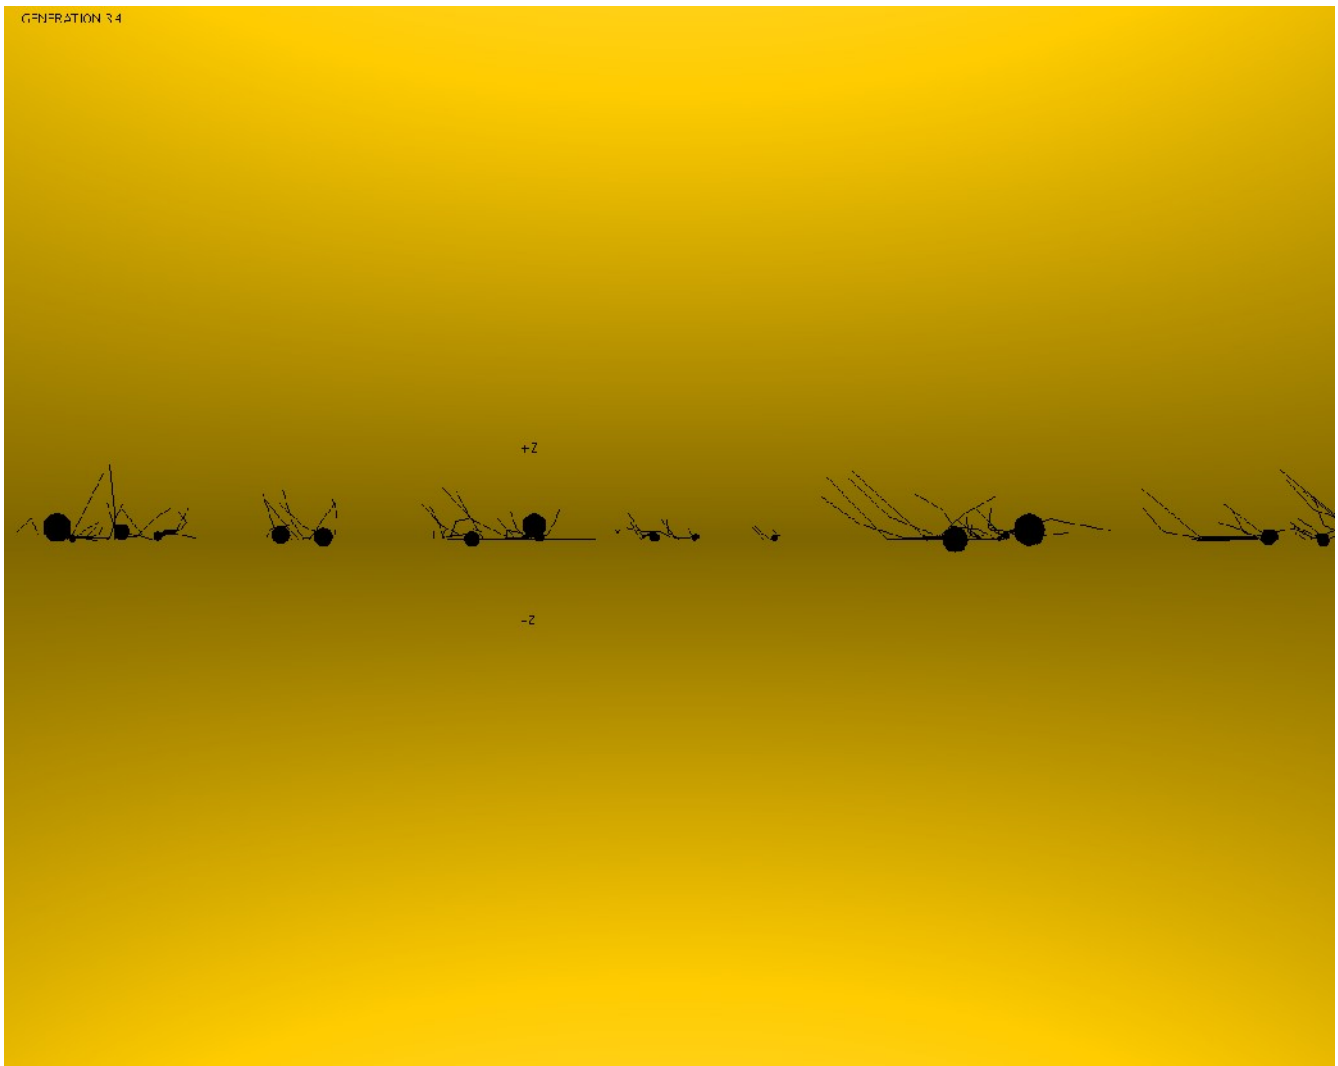


Figure 8: These stickles have learned to keep their legs out of the way, and slide along on their bellies while pulling themselves forward with their hands.

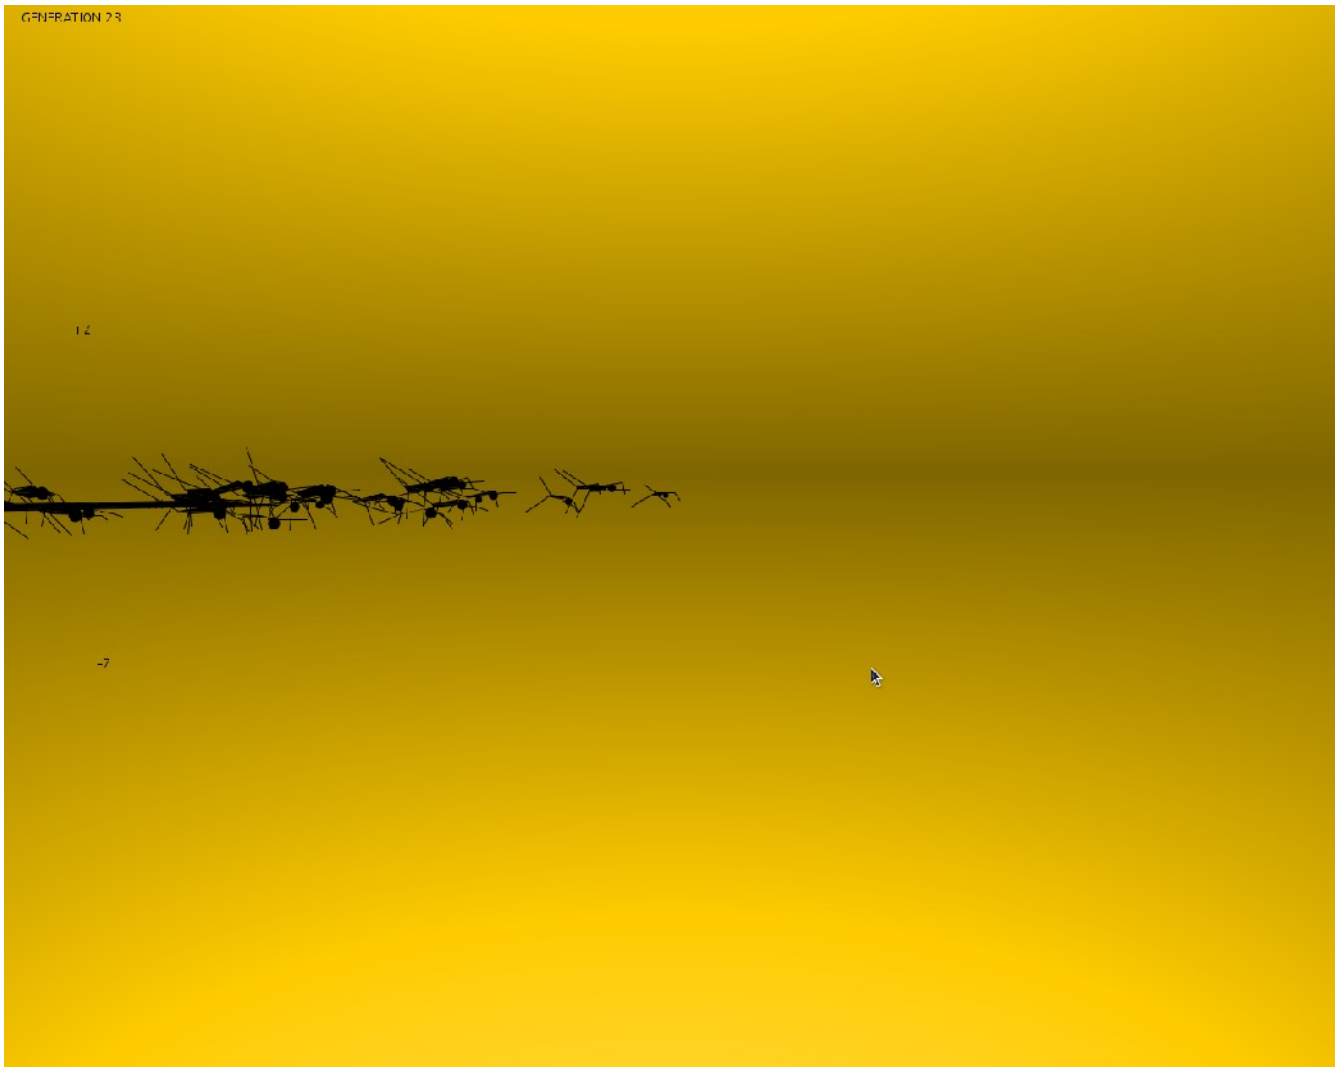




*Figure 9: A different angle of the previous stickles - their legs are also bent to the right. Some vertical tearing is present in this image: this is a side-effect of the screen-capture program used, and is meaningless.*



*Figure 10: These stickles pass their left leg outside then under them, and use it to push forward quickly. Meanwhile, their arms constantly crawl.*



*Figure 11: A low-angle view of the previous stickles, showing the leg-push moment.*



*Figure 12: Generation 35 of the same stickles for which the generations 0 to 2 are shown on pages 3 and 4. In addition to using their arms, these stickles use their heads to pull themselves along.*