

Comic Creation Game for StoryNet

Ryan Williams
breath@mit.edu

June 21, 2004

1 Motivation

The StoryNet project aims to collect information about stories to further machine understanding of the common sense aspect of human experience. A *story* by this definition is a causal series of events or states. The existing OpenMind website has been effective at gathering information about the relationships between everyday objects, such as “shampoo is found in the bathroom”. The Common Sense group also wishes to collect information about how events normally are sequenced, such as, “Eric was thirsty. Eric took milk out of the fridge. Eric drank the milk, and was satisfied.” Ministories such as these can teach a great deal about motivation and a combination of actions leading to a result.

Story knowledge will prove to be useful in a number of areas. Computer games could use the knowledge collected to auto-generate storylines and more realistic characters. Economists, statisticians, actuaries, and advertisers could use the story knowledge to model populations. The Cyc project uses a similar database of knowledge to check consistency across databases.

StoryNet will get some of its input the same way ConceptNet and LifeNet do: through mining the inputs of people on the Web. There is currently no official interface for this purpose. One interface exists that allows people to simply create stories in textual form, either by typing or by dragging and dropping prefabricated sentences, but it is not the ideal interface. I would like to make an interface that is so engaging that people keep on coming back and entering more and more sophisticated stories. This interface would be a game.

2 Game Overview

This game will be about comic creation. The most basic form of the game involves freeform comic creation. The player will create short comics by combining elements into panels, and adjusting a text description field for each panel. The game will exercise some degree of prediction to make the task of synchronizing the picture with the text easier.

The game’s playing field will consist of a workspace where the player builds their story, and one or more palettes from which they can select panels. The palettes will contain panel elements, already-filled panels, and series of multiple filled-in panels.

2.1 The panels

The panels are the fundamental building blocks of the comic. Each panel represents another state or action in the progression of the story. We don’t have the resources to create panels to represent every possible story action – the number of possibilities is as infinite as the number of stories. The panels will depict the action in a rather abstract and stylized way to get around this problem. The current plan is to create a finite set of a few dozen to a

few hundred panel parts that can be combined to represent a very wide variety of situations abstractly, and allow players to use their imaginations to fill in the missing visual gaps.

There will be a few types of objects associated with panels: a place where the panel is located, persons in the panel, objects, and actions being performed by persons or objects. Places can be represented by the background graphics, and people, objects, and actions by various colored shapes, as seen in Figure 1.

Figure 1: Sample workspace



Associated with each panel is its textual equivalent, one or more sentences. Since we represent the actors and scene with generic shapes, we rely on the players to appropriately label the objects; e.g. the green rectangle represents ice cream. I believe that it will be acceptable intellectually to players that the ice cream doesn't look realistic, as long as the green rectangle consistently is treated as ice cream. The game will help the player with this, by appropriately labeling the objects in the palette.

The primary player activity will be constructing and placing panels adjacent to each other. The player will pick a panel from a palette, or create a panel, and place it into sequence, and adjust the text if necessary to add more specificity than the graphics allow.

2.2 Causes & Effects

It is important for StoryNet to learn about causality in a more sophisticated way than simple sequentiality. In the ministory from Figure 1, there is a 'caused' graphical link that sits between panels two and three. The purpose of this graphic is to tell StoryNet that the action in the top frame directly *caused* the action in the bottom frame. Between the first and second panels there is a 'prerequisite' link between the two panels, indicating that Bob had to take the ice cream to the park in order to be able to give it to Sally. Other constructs that may be useful are 'preventative' links in which the action in one frame prevents something from happening, and 'unconnected' links that indicate unrelated frames.

The links look and act like puzzle pieces. In freeform mode the player may select any

links they wish, but the game may suggest relevant ones, and in some game modes the selection may be limited in order to increase the challenge.

We won't limit story construction to strictly sequential stories. It will be possible to construct parallel threads of causality within one story.

2.3 Autosuggestion

One of the goals of the graphical game is to engage players in a way that's not possible textually. Another goal is to encourage the players to expand the horizons of StoryNet's knowledge while keeping stories simple and direct enough that it's able to learn from them at all. Doing so involves a delicate balance of free-form input and constrained building-block construction. The approach that satisfies the two design goals of graphical interaction and limiting the input complexity is autosuggesting objects, actions, and text descriptions.

Autosuggestion builds off of the existing knowledge in ConceptNet, LifeNet, and the nascent StoryNet. Based on that knowledge and the contents of the panels adjacent to the current panel, the game will search for relevant actions and objects and add them to the palette for the player to use if necessary. In Figure 1, the objects 'cone' and 'tree' were suggested based on the presence of the ice cream and the panel's location at the park. The actions 'drop', 'taste', 'give' also represent autosuggestions. The player is still welcome to ignore the autosuggestions and manually create new objects and actions.

Autosuggestion will also be useful in preventing tedious text entry. The contents of each panel will suggest a set of possible descriptions. The game will attempt to use its existing ConceptNet and StoryNet knowledge to come up with a reasonable set of a few probable descriptive sentences. The player will be able to select or modify any of these, or write completely new text, perhaps through a drop-down menu. Giving the player building blocks in this way hopefully can achieve a balance between freedom and constraint, new knowledge and inordinately complex input.

3 Goals

One aspect of the game that isn't represented graphically in Figure 1 is goals. In order for a story to be meaningful, it must have obstacles being overcome on the path to a goal (stories do not differ too much from games in this respect). The StoryNet interface must be able to understand goals and how the actions in the story lead towards the goal. It's difficult to find a way to represent goals – some are abstract, some are concrete, different characters will have differing and conflicting goals, not to mention the fact that every character likely has a multiplicity of goals. One way to deal with this challenge in a way that respects the simple nature of StoryNet is to have goals represent concrete states; perhaps a specific panel will represent the 'goal'. In the Bob and Sally story, perhaps Bob's goal is to make Sally happy.

For freeform stories, the player will mark the goal panel somehow, at the same time identifying to which character the goal belongs. Each story may have many goal panels. An interesting extension may be to develop a mechanism for specifying subgoals, a task that may overlap with the case & effect links. This area remains under consideration.

4 Game Variants

The panel activity thus far described is a freeform play activity, and not a game. With a game, it's important to give the player sufficient challenge and direction without limiting

their freedom. One way to accomplish this is to constrain the player's choices for achieving a story goal. The player then has obstacles to overcome, and a quantifiable outcome of a completed story.

As described by Salen and Zimmerman, part of the definition of a game is that it has quantifiable outcomes. A program like this whose goal is the acquisition of knowledge cannot create quantifiable outcomes because it cannot judge the story the player creates. This is why the game modes are more about constraining the player's usable toolkit, because then the outcome becomes "did you use all the pieces or not?". Of course, players who wish only to win will write bad stories that don't make sense, and the game won't notice.

4.1 Fixed story points

The idea here is very similar to the OpenMind activity "Connect the sentences" in that the computer specifies the positions of two panels and their contents and requests that the player fill in the intermediate steps. The story might start out with a panel entitled "Bob is at the airport" and conclude with "Bob is drinking a cocktail on the beach", and the player's job is to insert the panels that complete the story.

4.2 Two-story connection

A more advanced version of the previous idea is that the computer specifies two or more ministories and asks the player to connect them. The player may be asked to connect them serially or in parallel.

4.3 Grouped panels

Here the goal is to create a cohesive story out of multiple ministory blocks of several panels each. The player would have to have some autonomy in selecting nouns or places so that the stories have more cohesion.

4.4 Time challenge

This subgame will be a timed puzzle game where the player races against the clock to use the pieces (a mixture of prefilled panels, and ministories, with a few blank panels to fill in otherwise intractable gaps) in a cohesive story.

5 Bootstrapping

Some of the more advanced game modes aren't possible without an existing body of knowledge about stories. If the computer is to string together panels to make a ministory for the player, it must already have some knowledge about how stories flow. Some of this knowledge already exists courtesy of the LifeNet and ConceptNet projects.

To some degree, the sophistication of the game itself and the appropriateness with which it can create or group panels will improve over the life of the project. As it grows more advanced, the knowledge base may be able to be used in game modes advanced beyond those described here. A mature knowledge base would be able to correct players who enter incongruous sequences. It might say, "the action in this panel doesn't follow from the action in the previous panel".

6 Player Persistence

Much like the OpenMind project interface, players should be able to access all of the stories that they have created. Unlike the OpenMind interface, players should be able to return to already-written stories and modify or add to them. The game can internally store the modified story as a separate entity from the original, or it may replace the original with the new one. The player should never lose their created stories.