Worlds Collide: Initial Developments in Ontologies for DOTA 2

Adam Houser, M.S. '15
Department of Industrial Engineering
University at Buffalo, State University of New York

0. ABSTRACT

The proliferation of high-speed Internet service, developments in computer graphics, and burgeoning interest in online multiplayer games have led to a rise in several well-known computer game titles that attract millions of players annually. However, new players are at a significant disadvantage to senior players, particularly where certain aspects of the game's ecosystem rely on specialized or nuanced knowledge. This is problematic for both the game's online community and software developer, as dissatisfied or frustrated new players who quit can lead to stagnant game growth and revenue. One potential solution may be the introduction of an ontology for a particular game, helping new players overcome steep learning curves and balance in-game experiences when matched against seasoned veterans. The author outlines developments in ontology creation for DOTA 2, a popular online multiplayer game, as well as directions for future development.

1. INTRODUCTION

One who develops ontologies about a given system or problem space seeks to formalize, in a fashion, knowledge of that system or problem space so that humans or machines can use them to accomplish desired tasks. This formalization of knowledge requires that ontology engineers annotate and categorize pieces of data within the problem space following standards agreed upon by all parties involved; because these standards almost always follow some sort of formal logic, the hope is that several different entities (computer systems, most generally, but humans as well) following the same logical templates will be able to use that problem space's data to draw conclusions about it.

One who plays video games likely does not care about ontological engineering. The gamer is concerned with the virtual world before them, learning that which is necessary for successful navigation of the virtual world: rules of gameplay, conditions for victory, different methods to attain that victory, tools that can be used within the virtual world, and a host of other features and factors. The only way for the gamer to take advantage of these features is to deduce and systematize them (often in mental models, although some gamers take the liberty of writing guides to games they have mastered). In other words, a game is a puzzle that a gamer attempts to solve by figuring out, organizing, and taking advantage of certain features of that puzzle.

However, the ontology engineer and the gamer may have more to learn from each other than seems apparent at first blush. Both are presented with a problem space; both problem spaces contain information critical to understanding the space; both seek to formalize their respective spaces. Ontology engineers and gamers also seek "solutions" to their problems spaces, although these solutions are a bit different: the engineer's puzzle is "solved" when the developed ontology enables successful information sharing and usage across systems, while the gamer's puzzle is "solved" once a set of victory conditions is met. Both solutions, though, are incomplete: the formalization of the ontology engineer's problem space (the ontology) can never be complete and perfect, just as the gamer's problem space (the game) can never be fully explored. Ontologies are never "completely finished" and video games are never "completely won."

These similarities indicate that a relationship between ontologies and video games may be fruitful, whereby an ontology may help gamers formalize their problem spaces and lead to puzzle solutions. The author contends that this very relation can be established between ontology and a video game titled Defense of the Ancients (DOTA) 2, a very popular online multiplayer game, by leveling the playing field for novice players through a formalization of the problem space.

2. BACKGROUND AND SCOPE

DOTA 2 is a team-based online multiplayer game that has been developed by the Valve Corporation. The game has a broad fan base in the United States and abroad, nominally hosting over 6.4 million unique players per month and nearly 500,000 concurrent players (that is, people who are playing at the same time). National and international tournaments regularly take place, resulting in prize purses valued at several hundred thousand dollars in cash; the DOTA 2 International 2013 competition sported a grand prize of $1.43 million to be split among a team of five players. Though interest in the game continues to rise around the world, new players can find themselves lost in a sea of rules and strategy that requires a careful balance of both force and nuance.

The premise of the game is simple: infiltrate the enemy base, destroy their primary structure, win the game. To do so, players on two teams (named "the Dire" and "the Radiant") select a unit, or "hero," from a list of 103 available. Players can see the position of other allied heroes but have control only over his or her individual hero, moving it across the map either alone or with teammates to achieve certain objectives. These objectives are the same for every game played.

Figure 1, below, shows the worldmap of the DOTA 2 game. Several things are important to note. Heroes of the Radiant (green) start bottom-left and work towards top-right, while the Dire (red) proceed as the mirror opposite. There are three paths through the forest, called lanes, which are designated as "top," "mid" or middle, and "bot" or bottom. These lanes sport guard towers that engage enemies and support allies that come within range; boxes indicate their locations on the map, with color indicating ownership of the towers. To achieve victory, a team must destroy the guard towers in at least one entire lane, moving in towards the enemy's base and destroying the towers in sequential order. Once these towers have been destroyed, the team can attack the enemy's Ancient, the primary structure at hub of each base in the extreme upper-right and lower-left corners, with the intent of destroying it. Once the enemy Ancient is eliminated, the game is finished.

Figure 1. The DOTA 2 worldmap.

Because things are never as easy as they seem, however, several confounding factors enter play. First is the enemy team, composed of five heroes with the ability to kill opponents and wreak an equal amount of havoc on the opposing team. Pushes into enemy territory must be heavily weighed against the specter of counterattack and team-vs.-single hero assassinations, called "ganks," which can result in swift (perhaps catastrophic) shifts in game momentum. Second is a dynamic map that provides alternative routes, hiding spots, and quick paths through the jungle that allow players to gang up on and surprise enemy heroes who may be caught unaware or out in the open, away from the protection of towers and allies. Third is an automatically regenerating army of weak minions, called "creeps," that spawns in waves and move towards the enemy base in fixed paths down each of the lanes simultaneously. These are only controlled by the computer and exhibit no real strategy beyond attacking whatever enemies appear within range. Rather, the ostensible purpose of creeps is to provide fodder for heroes, who can gain experience points and gold coins for delivering the fatal blow to enemy creeps. Heroes that accumulate experience points gain power, making them tougher, faster, stronger, and unlocking special abilities that they can use in combat against the opposing team. Gold coins allow heroes to

purchase items at shops scattered around the map, further bolstering heroes and—depending on the items selected—giving heroes new abilities that can swing the tide of war in one direction or another with ease.

The scope of this ontology will be broad rather than deep, focusing on demarcating knowledge of "tangibles" (heroes, structures, neutral units, and the item purchasing and creation process) as well as "incorporeals" (hero actions and conditions) native to DOTA 2, with the ontology's greatest detail and development being placed upon the item creation process. It has been designed within the framework of the Basic Formal Ontology (BFO) v1.1, following its superclass/subclass distinctions and relationships. Annotations can be found throughout the ontology explaining the definitions of terms, how they are expressed in RDF "triple" format, and their relation to the SNAP and SPAN differentiation within BFO. Because an ontology is most useful if it can be integrated with other ontologies, it was the author's intention to construct the DOTA 2 ontology within the BFO framework to facilitate this eventual integration. With this format in mind, we should now turn to the major structure of the ontology and explanations of these constructions.[1]

3. GENERAL NOTES ON STRUCTURE

Annotating the problem space through BFO helps to establish features of the problem space in an easy to understand, easy to describe systematic format. The ontology in its present construction makes heavy use of independent continuants. For the ontology to work properly, one has to reason as if the digital entities seen and used within DOTA 2 actually do exist and are not merely ephemeral digital constructs that disappear if the game is not being played. Once this has been established, these digital entities can be considered as belonging to one of four classes: items, heroes, creeps, and structures. Items are tools or weapons used in the game by heroes, digital persons or creatures that humans can play and control within the game. Heroes can use items to dispatch other enemy heroes or creeps, AI-controlled minor creatures present throughout the map. Heroes can also destroy structures, and in fact destroying certain structures is the only way to achieve victory within the game. An early problem with constructing the ontology was a decision to categorize structures as an independent continuant or a spatial region, one of BFO's three classes of continuants. It was decided to classify the game's structures as independent continuants rather than spatial regions because DOTA 2 places an emphasis on the structure's existence (or absence, after it has been destroyed) as a condition for victory; in other words, while they do occupy spatial regions on the map they are seen particularly as unique digital entities with an emphasis on existence rather than location (because all structures exist in exactly the same location for every round of DOTA 2 played by every player). That structures occupy space on the map is secondary to whether it exists, leading to its classification as an independent continuant.

What, then, should be classified as a spatial region if not structures that occupy space? The answer seems to be spatial regions whose existence relate in no way to victory conditions and are not "interactive" (in other words, they cannot be destroyed as structures can). Contained herein are four major spatial regions with importance to players: bases, lanes, the jungle, and the river, each with specific subclasses expressed where they are necessary.

A final piece of the DOTA 2 ontology deserving of focus is the occurrent, or processes that have temporal components. In the current construction, this ontology describes both character states and game states. These are not roles or attributes of heroes (which may be more suitable for description elsewhere in the ontology, such as within the dependent continuant class), but are rather processes in which heroes can be engaged. For example, a hero moving down a lane with allied creeps with the intent of destroying the towers in that lane is "pushing" the lane. Pushing is occurring, but it cannot occur without an independent continuant (the hero) to do so, so it becomes a processural occurrent. The same can be said of victory and defeat, game states that exist once a game has been won or lost and is dependent on the success or failure of a team of heroes. These are not character states, because each hero does not

[1] Please see Appendix A for images of the ontology visualized in Protégé.

win or lose; rather, the match in its entirety is either won or lost and is therefore a processural occurrent.

4. ONTOLOGY EVALUATION

Testing ontologies incorporate a number of different methods, some of which include logical parsing and others that include knowledge store querying through techniques like SPARQL. Because this ontology is an initial attempt at annotating the problem space of DOTA 2, attempts at performing something akin to SPARQL querying will undoubtedly return incomplete results that do not properly showcase the structure of the ontology in its present form. Logical parsing could be more effective here, and a good way to do so would be to examine the ontology through application of the "true path" rule.

This rule describes a technique in which the path from the top to the bottom of an ontology's hierarchy is examined, with the goal of seeing that superclass/subclass distinctions hold true for each class examined and that no problems of multiple inheritance are found (or, at least, multiple inheritance instances are kept to an absolute minimum). To demonstrate the true path rule in action, consider the following two classes: Ring_of_Health and Dire_Base.

The Ring_of_Health is an item that can be purchased by heroes within the game. The Ring_of_Health is a secret_item, which is an item, which is a digital_entity, which is an independent_continuant, which is a continuant, which is an entity, which is a Thing. All of these edges represent sound relations between nodes, and by separating the items into their three respective classifications the true path rule is maintained. In other words, Ring_of_Health cannot be both a secret_item and a unique_item, nor can it be a spatial_region or anything else. The path from Ring_of_Health to Thing follows one strict path and is traversable both forward and backward. This also holds for Dire_Base: it is a Base, which is a spatial_region, which is a digital_entity, which is an independent continuant, which is a continuant, which is an entity, which is a Thing. Here the previous decision to maintain spatial regions separate from structures pays off. Dire_Base is not a structure but a region of the map where heroes can return for safety; certain benefits are conferred to heroes who are within a certain radius of their respective bases. If this distinction were not made early on, then it may be possible to violate the true path rule by expressing a base both in terms of the structures that are important to it (the Fountain, for example, is a structure located within a spatial region but is an independent continuant that should be thought of in terms of its existence, not whether it occupies space) and the area it occupies. Because one can traverse this path up and down the ontology while keeping this distinction in mind, the true path rule is in no danger of violation.

5. COMPARISON TO EXISTING WORK

Evaluation of the ontology is important for assessing its value and relationship to existing work; as was mentioned before, an ontology is most useful if it can be integrated with other ontologies. The chances of the DOTA 2 ontology's success and regard as a serious attempt at ontologizing a video game's problem space depends on the quality of integration. As such, this ontology has been developed in accordance with BFO v1.1 and the Information Artifact Ontology, or IAO. Efforts to design ontologies as interoperable and integrative emphasize the usefulness and prevalence of these two ontologies, and because of this the author thought it wise to also construct the DOTA 2 ontology in the same spirit.

Beyond these two systems, little other serious work on video game ontologies exists in published format on the Internet. Search engine queries for "video game ontologies" or "video game ontology" yielded at most 68,800 results, but a vast majority of these hits are either junk returns (websites with aggregates of words, stored to garner results or act like web crawler honeypots) or are related to ontologies of games that are descriptions of terms rather than workable hierarchies and searchable knowledge stores. The only significant return located during Internet searching was a video game ontology found on a site called BitBucket and written by Mike Cutalo, Todd Conway, and Andrew Bashore.[2] The ontology is one of video games in a general sense, including many different types of video games, game platforms, and some

[2] https://bitbucket.org/bashtech/video-game-ontology/overview

terminology associated with their creation and play. However, when the .owl file is loaded (see Figure 2, below) it becomes apparent that it does not follow BFO v1.1 at all and introduces no real hierarchy that could later be expanded or modified. This is rather problematic and indicates that it would not be a good candidate for integration into this ontology.

Figure 2. Detail of the Video Game Ontology

Furthermore, there seem to be some violations of the true path rule. Consider the class "Vehicle Video Game," which is categorized under the Simulation Video Game superclass. It is arguable that a vehicle video game could also be classified as a racing video game, meaning that there could be confusion during database querying and that there is not one true path that encompasses these classes of racing video games. Another possible violation in the existence of "Dungeon Crawler Video Game" as separate from "Role Playing Video Game," both of which are separate from "Adventure Video Game." The author speaks from years of experience when he says that dungeon crawlers (video games in which the player controls a character, much like DOTA 2's heroes, but instead of playing on a large outdoor map the hero progresses through a dungeon fighting monsters and solving puzzles) are usually role playing video games, and that a role playing video game is usually considered an adventure video game. Because of these and other problems, the author therefore decided against integration (barring extensive work to clean up this Video Game Ontology) in favor of creating a separate ontology that followed existing structure guidelines more closely and provided easy access and modularity for future work.

6. POTENTIAL APPLICATIONS

The goal of this and other ontologies is to systematically categorize and annotate information found within a problem space, allowing users to query that information and extract desired information. The DOTA 2 ontology is designed to do just that, allowing users to query the ontology to extract information about the game. As the ontology is improved, the quality and integration of this information will yield better query results and an increase in utility among those using the ontology to learn something about the game.

One of the major applications is the reduction in cognitive load among novice players. Every game is codified through rules and information about that game, and the difficulty of learning and remembering this information is a function of several different factors (the goals of the game, victory conditions, the tools and implements used, et cetera). The initial stages of game playing can be the most difficult as the novice player seeks to integrate newly learned information into a survey of the game's problem space, and this integration of information can result in frustration for the novice as they seek to apply this information to a new problem space. Online gaming is the same, particularly where playing against other humans in remote locations is concerned. It may be difficult to assess the skill level of the other player or players before the game begins, and frustration can result in leaving the game and having a generally negative experience. Loss of new players is something that no game publisher desires, and safeguards have been erected to prevent this from happening (consider the match-making software most online games now employ, whereby players of similar skill and experience are matched up to compete against each other). However, this does not necessarily work and does not help the novice player learn and apply the rules of the game, so additional actions could be taken. Enter the ontology.

One particular area within DOTA 2 in which the ontology could be useful is the item creation process. Found under independent continuants in the author's ontology, items are tools and weapons that heroes can purchase and construct within the game to bring them increased power and ability. Teams that are most often victorious give special consideration to the items their heroes possess. The problem for the novice player is that the item creation process is complicated and laden with information: some items combine in certain ways to create more powerful items; many items come with restrictions on use or utility; the resource pool of the hero is finite, meaning that a player must be very careful with how the he or she progresses through the item creation process. In all, 128 items are available for purchase or creation, meaning that a player must work through this list to determine exactly what is to be purchased or created, how long it will take to do so, and what benefits or penalties (if any) it will confer on the hero or the team, in a very rapid fashion without stopping game play. This is a task any novice could find daunting and is not helped by match making software or anything else that is present within the game when one begins to play.

An ontology would be a great boon for novice players seeking to make this faster, smarter, and more efficient. Once the ontology has annotated and described the details of the item creation process, a novice could submit a series of queries to the database to determine what types of items should be created, how much they would cost, and what types of benefits they would confer. For example, consider Aghanim's Scepter, an upgrade item that's available to be created within DOTA 2. Detail of its .owl file entry can be found below in Figure 3. This item is an upgrade item that is comprised of four different lesser items: Blade of Alacrity, Ogre Club, Staff of Wizardry, and Point Booster; in other words, when a player combines these four items together they receive Aghanim's Scepter. Figure 4 then shows what the reasoner would show upon successful execution: that the Blade of Alacrity is a part of Aghanim's Scepter. These items would all then have monetary costs, presented in gold, affixed as data properties.

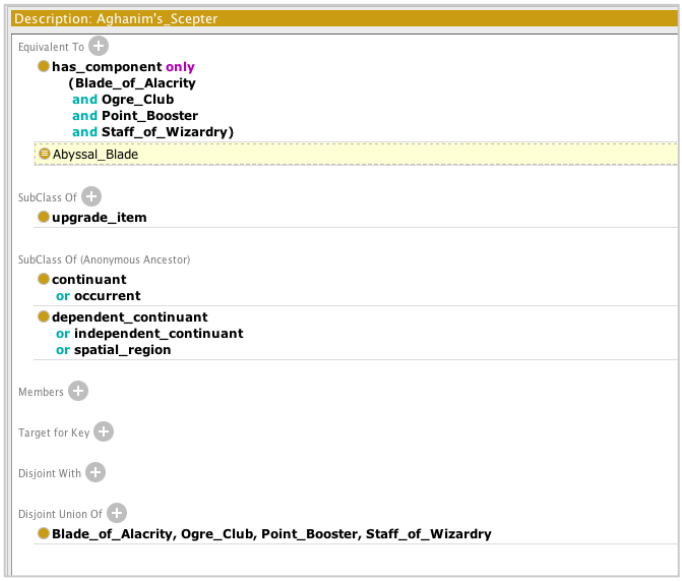

Figure 3. Detail from Aghanim's Scepter.

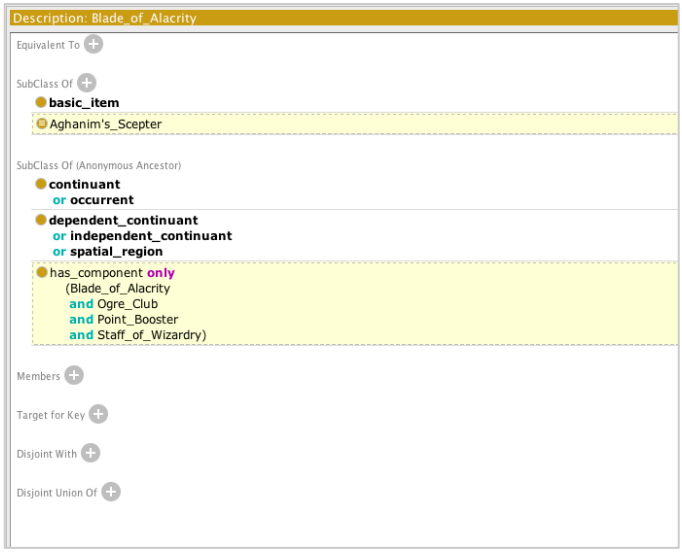

Figure 4. Detail from Blade of Alacrity.

Once these relations and data properties have been properly annotated and expressed, a novice player using SPARQL could enter a number of different things to determine what to buy or how much gold to save up. For example, a player could establish that he or she possesses Ogre Club and Blade of Alacrity, but need to determine what else they could make from these individual items; alternatively, a player could express that they have a Blade of Alacrity and 3,600 gold saved up from completing different challenges in the game, and he or she needs to know what is available for purchase and would be synergistic with the Blade. These queries would all be possible with the DOTA 2 ontology.

A potential problem is that novice players in the middle of a game do not have the time to run SPARQL queries using something like DBpedia. A solution would be to develop a GUI with drop-down menus, text fields, or perhaps drag-and-drop menu items to allow users to enter their gold amounts, drag pictures of what items their heroes currently own, and select features from menus regarding what they are looking for in items available for purchase. The software's front end could use SPARQL on the back end to query the ontology, return results, and parse them into graphical or textual information that novices can use with ease. It could be a quick, easy, and effective way for inexperienced players to decide which items are needed and what in-game actions they should take to secure said items. Given time these novice players will still learn the intricacies of the item creation process, but it will help new players immediately begin competing and hopefully alleviate some of the frustration that may come with the learning curve.

It should be mentioned that this is but one potential application of this ontology. As development continues, it may be possible to annotate DOTA 2 heroes (103 in total and already added to the ontology within the Digital Entity superclass) with their primary traits (Dexterity, Intelligence, or Strength) and special abilities, then use the same GUI and database querying system to find the items available for purchase that are most helpful to the particular hero being played. One could also query the database to figure out which heroes or items are useful for different actions within the game, such as farming or jungling. The possibilities with a well-annotated data set constructed in an ontology are vast, and it is the opinion of the author that this could be incredibly valuable to novice players as construction on the ontology continues.

7. FUTURE WORK

Because an ontology is never complete, there will always be work to do and information to add. There are several things that can be done right now to further develop this ontology.

- Finalize the item creation hierarchies by adding in all of the item relations. This is a particularly vexing problem with the ontology as it currently stands. The reasoner does not like how the ontology is parsed and will not automatically populate certain fields of data, but it is not exactly clear why. For example, when lesser items are mentioned as components of greater items as shown in Figures 3 and 4, the subclass expressions are not specific to greater items that contain that lesser item. So if one continues by adding more lesser item recipes, they all end up being displayed as subclasses of every other lesser item. In other words, clicking on a lesser item shows that it is the subclass of EVERY other greater item, regardless of whether it is or is not contained within the other greater items. This is something that should be fixed if the SPARQL query process is to be most effective.
- Add information to further develop heroes. After items, heroes are the second most complex digital entities within DOTA 2. They are arguably more nuanced and so provide a great amount of data to catalog for the database. If one could add this information (some of which is mentioned above), one could turn this ontology into a very powerful tool for synergistic item and hero selection and creation.
- Give thought and annotate strategy. Most of this ontology's development has gone into tangible entities, with not much having been put into incorporeals such a strategic moves. Their organization will require a great deal of thought but is entirely possible within the constraints introduced within this ontology.

It is the hope of the author that this ontology makes its way to someone who can carry on this work and perhaps help integrate the information sets found within DOTA 2. Such a program could be an extraordinary help to novice players if it could be made into a working prototype. Hopefully the work done thus far on this ontology is enough to get us started.

APPENDIX A: The DOTA 2 Ontology

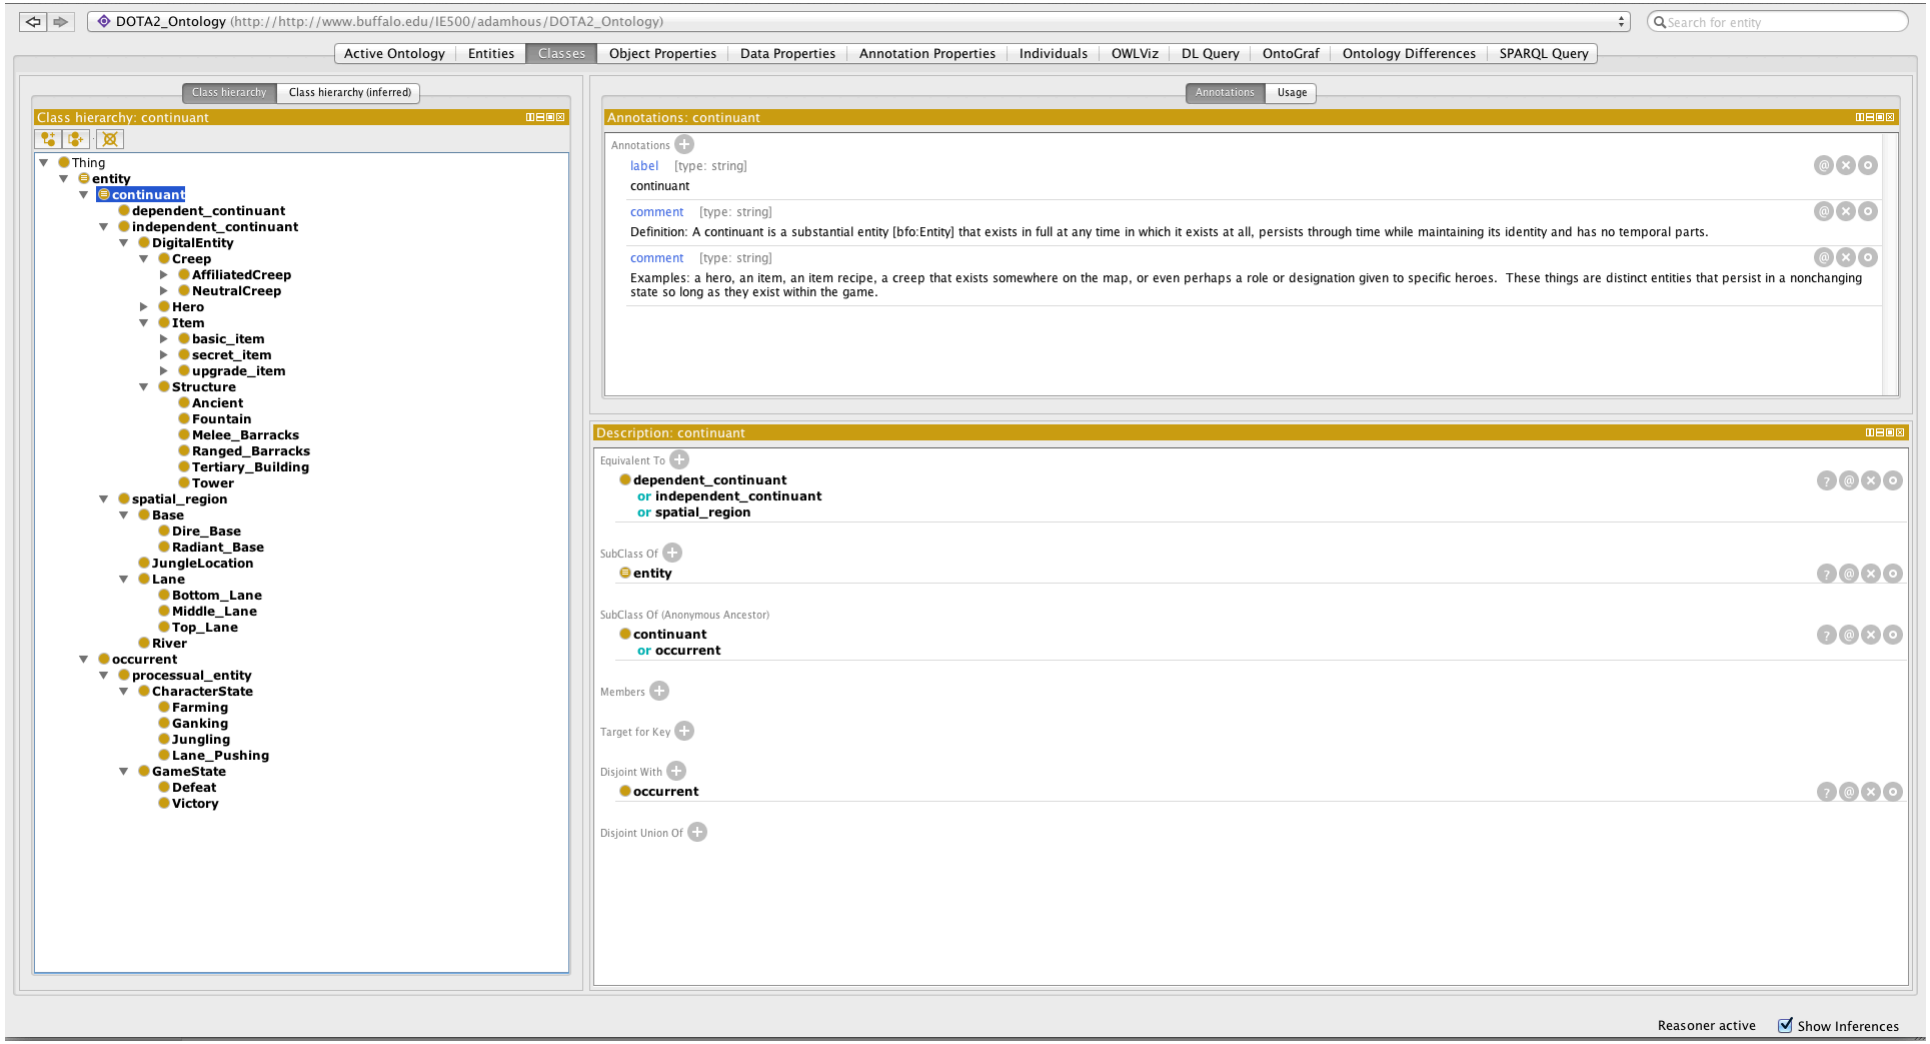

APPENDIX B: The DOTA 2 Ontology, Visualized in OntoGraf