Envisioning With Weblogs

Andrew S. Gordon
(University of Southern California, Los Angeles, CA USA
gordon@ict.usc.edu)

Reid Swanson
(University of Southern California, Los Angeles, CA USA
swansonr@ict.usc.edu)

Abstract: In this position paper we present a vision of how the stories that people tell in Internet weblogs can be used directly for automated commonsense reasoning, specifically to support the core envisionment functions of event prediction, explanation, and imagination.

1 Commonsense Knowledge From The Social Web

Commonsense reasoning was identified as a key problem in the field of Artificial Intelligence at its inception fifty years ago. Since then, the predominant approach to this problem has involved the logical formalization of commonsense knowledge. Although significant investment has been directed toward the authoring of formal content theories by professional knowledge engineers and researchers, only slow progress has been made toward the central goal of achieving broad-coverage commonsense formal knowledge resources [Davis, 2004]. The beginning of the current decade saw a shift away from the predominant approach, with the labor of commonsense knowledge engineering being distributed across members of large volunteer communities. Exemplified by the Open Mind Common Sense project [Singh, 2004], this work differed from the previous approach in three ways. First, natural language was used to represent commonsense knowledge, rather than logical notation. Second, knowledge bases were created as a product of contributions from thousands of (untrained) web volunteers. Third, new methods of automated reasoning were used to support commonsense inference in applications, instead of logically sound theorem-proving techniques. By pursuing a pragmatic approach to the problem of scale, this line of research has shown some success in integrating commonsense reasoning into real-world applications [Lieberman, 2004].

While the use of large volunteer communities is the most prominent characteristic of the approach, it is the use of natural language representations and alternative forms of inference that is now fuelling another shift in the way that commonsense reasoning research is pursued. The thousands of volunteers who contribute natural language to Open Mind Common Sense pale in comparison to the tens of millions of people who contribute natural language into their weblogs, online forums, and social networking websites. The Social Web has more breadth and depth of commonsense knowledge

than will ever be seen in hand-authored logical formalizations or in volunteer-contributed knowledge bases. This observation has recently spurred interest in commonsense information extraction from the web, where the *scruffy* text of social media is mined and repackaged into *neat* formalisms that can be used in traditional styles of automated reasoning. The fallacy of this approach is that using neat knowledge bases for general-purpose commonsense reasoning remains an unsolved problem, plagued by brittleness, the challenges of non-monotonic and default inference, and intractably large search spaces. Instead of squeezing the Social Web into the straightjacket of traditional approaches, we should be exploring ways to utilize its unique characteristics more directly in support of commonsense reasoning.

2   Envisioning With Weblogs

Reasoning about events and their causal and temporal relationships is the most central of all human cognitive abilities. Predicting the consequences of events in the world enables us to form expectations, identify potential threats to our goals, and consider alternative courses of action in the development of plans. Imagining the antecedent causes of events allows us to generate explanations for our expectation failures, and for the failure of our plans. Imagining events outside of our own experience enables our minds to wander into daydreams, consider the narratives of other people's experiences, and enjoy fictional stories. Collectively, these processes of event prediction, explanation, and imagination can be described using the umbrella term *envisioning*. Human-like Artificial Intelligence requires the development of robust general-purpose algorithms and formalisms for envisioning, and this has been the focus of most previous work in this area. However, human-like envisioning also requires a human-like model of how the world works, i.e. commonsense knowledge about the world in which people live. The lack of commonsense knowledge is one of the central reasons that we do not have successful envisioning systems today, and is subsequently a reason that we lack human-like artificial intelligence with the abilities to make explanations, draw expectations, devise plans in the real world, anticipate real-world threats, understand stories, and generate creative narratives.

   Internet weblogs contain massive amounts of information about causal and temporal relationships between events in the world, contributed by tens of millions of authors. Much of this event information takes the form of first-person narratives of people's past experiences. The following (real) passage is typical of the stories found in weblog entries:

> "*I cracked the egg into the bowl and then I saw it, yes a baby chicken was in the egg that was going to be our breakfast. I felt like I might be sick, but the rest of my family found this to be very interesting.*"

   This passage is particularly interesting from a research perspective because of a historical interest in the "egg-cracking problem," which served as a non-trivial benchmark problem in commonsense physical reasoning [Shanahan, 2004]. In its original form, the problem is to predict what would happen if the standard culinary practice of cracking an egg against a mixing bowl were changed in specific ways: What if the egg hits the edge of the bowl too quickly? What if the bowl is upside-

down? What if the egg is hard-boiled? Without resorting to logical formalization, this weblog author has just told us (in English) what happens when there is a live chicken inside the egg, namely that people will either feel like they might become sick or find it very interesting. In addition, the author has let us know that eggs turn into breakfast. Equally good answers for the original variants of the egg-cracking problem are readily available in weblog stories, with the exception of the most esoteric versions of the problem. The challenge we face is building new reasoning systems that can directly use this textual information in support of the basic processes of envisioning.

The genre of the everyday weblog story has both strengths and weaknesses as a source of commonsense knowledge, requiring that special consideration of the characteristics of this genre be taken into account. Schank & Abelson [Schank, 1995] have long argued that conversational stories of this sort are distinguished by the way that they violate commonsense expectations, suggesting that the knowledge extracted directly from stories may be the exact opposite of what is needed. This attribute is reflected in the egg-cracking passage above; we do not want to conclude from this text that baby chickens are expected to fall out of eggs that we crack. More subtly, there are a number of incorrect inferences in this passage that are more difficult to avoid: should we infer from this passage that people do not eat chickens? Problems of this sort have previously been used as a justification for the investing in the creation of hand-authored commonsense content theories. However, the massive scale of available content affords different approaches, e.g. the application of statistical, evidence-based techniques. To pursue these new approaches, however, the focus of research in the area of commonsense reasoning needs to shift toward a new set of research problems: story acquisition, story analysis, and story-based envisioning.

First, new story acquisition technologies are needed to automatically collect massive amounts of stories. Gordon et al. [Gordon, 2007] developed automatic techniques for extracting stories from weblog entries, where only 17% of the text is story-like in character. Using a machine learning text classification approach, this work obtained reasonable levels of precision and recall, but still much less than the inter-rater agreement among annotators of the training data. The further development of fast and accurate story acquisition techniques would allow for the creation and maintenance of extremely large corpora of textual stories from weblogs and other sources in the Social Web. The development of gold-standard training and evaluation corpora with high levels of inter-rater agreement would facilitate progress in this area.

Second, some amount of syntactic and semantic processing of the natural language text of stories will be necessary, particularly to identify events within these stories and the causal and temporal relationships between them. For some envisioning tasks, it may be reasonable to simply assume that sequences of story sentences map directly onto ordered sequences of events in the world. For envisioning that requires more precision, it may be advantageous to explicitly delineate events in narrative text and identify the semantic relationships between them. Recent work in this area has focused on machine learning approaches using annotated training corpora, e.g. TimeML annotations of news articles [Mani, 2006]. The utility of this technology in support of commonsense reasoning could be improved by shifting attention away from news article text and toward the stories of people's experiences, as can be found on the Social Web. Many of the existing event-analysis technologies would be

suitable for this genre, but would require the development of new training and evaluation corpora of annotated story text.

Third, new methods are needed to use the event sequences identified in stories to support automated commonsense reasoning. The basic reasoning tasks of envisioning (prediction, explanation, imagining) each require the capacity to produce a new event with a specific semantic relationship, given a state or event description as input. One strategy is to learn a model of event transitions using stories on the Social Web as training examples. This approach is exemplified in the work of Manshadi et al. [Manshadi, 2008], where language-modelling techniques were used to learn a probabilistic model of event sequences using stories extracted from weblogs. While this line of research is promising, we believe that alternative approaches to envisioning should also be explored. In particular, case-based reasoning techniques seem particularly well suited to this task, where event sequences serve as cases in an episodic memory, and the processes of envisioning are recast as textual retrieval and adaptation problems.

Acknowledgements

References

[Davis, 2004] Davis, E. & Morgenstern, L. (2004) Introduction: Progress in formal commonsense reasoning, Artificial Intelligence 153(1-2):1-12.

[Gordon, 2007] Gordon, A., Cao, Q., & Swanson, R. (2007) Automated Story Capture From Internet Weblogs. 4th Intl. Conference on Knowledge Capture, Whistler, BC.

[Lieberman, 2004] Lieberman, H., Liu, H., Singh, P., and Barry, B. (2004). Beating common sense into interactive applications. AI Magazine 25(4):63-76.

[Mani, 2006] Mani, I., Verhagen, M., Wellner, B., Lee C., & Pustejovsky, J. (2006) Machine Learning of Temporal Relations. 46th Meeting of the Association for Computational Linguistics, Sydney, Australia.

[Manshadi, 2008] Manshadi, M., Swanson, R., & Gordon, A. (2008) Learning a Probabilistic Model of Event Sequences From Internet Weblog Stories. 21st International FLAIRS Conference, Coconut Grove, FL.

[Schank, 1995] Schank, R., & Abelson, R. (1995). Knowledge and memory: The real story. In Wyer (Ed.), Knowledge and memory: The real story. Mahwah, NJ: Lawrence Erlbaum.

[Shanahan, 2004] Shanahan, M. (2004). An attempt to formalise a non-trivial benchmark problem in common sense reasoning. Artificial Intelligence 153(1-2):141-165.

[Singh, 2004] Singh, P., Barry, B., & Liu, H. (2004). Teaching machines about everyday life. BT Technology Journal 22(4):227-240.