Bursting the Filter Bubble:

From Echo Chambers to Speech Balloons with Data-Driven Comics

Tony Veale
University College Dublin (UCD)
Dublin, Ireland
tony.veale@UCD.ie

***Abstract* — Every expressive medium allows us to ground meaning in different ways. Comics, or the so-called 9th art (after film and TV), are sequential integrations of words and images that offer more possibilities than either words or images alone can offer. Like any art form, comics settle into clichéd norms – panels, balloons, tails – and give rise to genres that dominate at the expense of all others, such as the superhero genre. Yet comics can also use their vivid emotional expressivity to imbue physical actions with potent feelings and to lend intuitive immediacy to the gamut of human concerns. This paper considers comics as a medium for storytelling and as a data-driven content representation from which machines can shape specific meanings for particular audiences. In particular, we aim to generate engaging and unthreatening content that offers a balanced view of highly contentious topics of current debate. In doing so, we seek to foster greater communication between siloed users of social media who often talk about, but rarely talk to, those they most disagree with, so as to break down the echo chambers that work to keep these disputing parties apart from each other.**

Keywords — comics, story-generation, echo chambers, Twitter

Introduction

Online debate around divisive topics has become increasingly fractured, leading to the emergence of ``echo chambers'' in which disputants communicate almost exclusively with those who hold compatible views. Moreover, social media often makes it harder to disagree without also becoming disagreeable. To inhibit the growth of echo chambers and to expose disputants to both sides of an argument – in ways that encourage dialogue across the divide – we aim to automate the generation of creative interventions into otherwise insular online debates. On highly echoic platforms such as Twitter, bot-driven interventions run contrary to best practices, and may be reported as an abuse of the system. However, passive interventions can instead use story generation to dramatize an ongoing debate. If the stories so generated are engaging and balanced, and are aptly labeled with attested hashtags, they can draw users to a bot's content, thus avoiding any need for a bot to elbow its content into a live conversation. The *Excelsior* system, as described here, aims for amusing, even-handed engagement by packaging its data-driven stories as comic strips which integrate two sides of any argument into a single visual intervention.

Although frequently packaged in a disposable form, comics have been described as a sequential art by Eisner [1] – for whom the Eisner prize in comics is named – and as the ninth art by Maurice De Bevere, the creator of Lucky Luke. At its simplest, a comic strip is a sequence of framed snap shots, called panels, separated by thin bands of whitespace, called gutters. Each panel is usually square or rectangular, typically framed in a black border, and sometimes labeled with a caption above or below the scene depicted within. A typical panel contains a mix of textual and visual elements, to depict a specific action in a certain setting, and to record any words that are spoken (or thought) in context. Those words are most often contained within text balloons, either rounded speech balloons or fluffy, cloud-like thought balloons, whose tails link them to the vocalizing characters.

These conventions have become the stuff of cliché, but as McCloud [2] has shown, this normative grammar of comics allows for a great deal of divergence and creativity. Indeed, even text balloons can vary tremendously from one context to another, to shape meaning as well as just contain it [3]. Although the ease with which children adapt to the medium's mechanisms allows some to dismiss it as juvenile, this ease also reflects the depth of the cognitive foundations in which the medium is rooted [4,5]. For instance, one intuitively reads a comic strip in the order one reads a text, from top to bottom and left to right in the West, and from right to left and back to front in the East. No one needs to be taught how to read a comic strip. Readers simply adapt to the blended medium as a new form of visual reading.

Although comics are entertaining in their own right, we explore a practical use of the medium here. Consider why they are called 'comics' or 'funny-books' in the first place: the name is a carry over from the earliest newspaper strips in which short, whimsical diversions were illustrated; the first American comic book, *Famous Funnies*, repackaged these newspaper funny pages as a standalone periodical. Even serious comic books – which some now call *Graphic Novels* – still carry traces of the comical and unserious. The current work exploits these vestiges to package polarizing and unwelcome meanings in less disagreeable forms. Those meanings arise in heated online debates, such as the Twitter (now X) debate on vaccines, in which disputants show a baleful tendency to dig in, tune out and listen only to those on the same side. Machines can help to break down these "echo chambers" [6] by making targeted interventions into the debate, using comics to summarize and distill the main arguments on both sides.

We shall present two very different approaches to the problem, reflecting the dramatic changes in AI and NLP that have reshaped the field in recent months. The first is symbolic and rule-driven; the second exploits Large Language Models (LLMs) to do much of the work of the many different modules of the symbolic approach, albeit with greater fluency, insight and creative diversity. However, each relies on the same underlying comics representation in XML, and the same rendering mechanisms, to produce their end results. Although the symbolic system is rigid and stilted in its output, it is still capable of greater topicality and responsiveness than the LLM. We thus see obvious benefits to both approaches, and are seeking to synthesize their comparative advantages in a single framework.

References

[1] Eisner, W. (1985). Comics & Sequential Art. Tamarac, Florida: Poorhouse Press.

[2] McCloud, S. (1993). Understanding Comics: The Invisible Art. New York: Harper Collins.

[3] Forceville, C., Veale, T., & Feyaerts, K. (2010). Balloonics: The visuals of balloons in comics. The rise and reason of comics and graphic literature: Critical essays on the form. J. Goggin & D. Hassler-Forest (Eds.). Jefferson NC: McFarland, pp 56–73.

[4] Cohn, N.. (2013). The Visual Language of Comics: Introduction to the Structure and Cognition of Sequential Images. London: Bloomsbury

[5] Cohn, N. & Magliano, J.P. (2020). Visual Narrative Research: An Emerging Field in Cognitive Science. Topics in Cognitive Science, 12:197–223.

[6] Barberá, P.; Jost, J.; Nagler, J.; and Tucker, J. 2015. Tweeting from left to right: Is online political communication more than an echo chamber? Psychological Science 26(10):1531–1542.