

The Funhouse Mirror Has Two Sides: Visual Storification of Debates with Comics

Tony Veale¹

¹*School of Computer Science, University College Dublin, Ireland*

Abstract

Online debate around controversial topics has become increasingly fractured and entrenched, leading to the formation of *echo chambers* in which disputants communicate chiefly with those who hold compatible views. To inhibit the growth of such chambers, and to expose disputants to both sides of an argument – perhaps in ways that lead them to argue directly with their opposing counterparts – we can automate the generation of “interventions” into otherwise insular online discussions. However, on highly echoic platforms such as Twitter, direct interventions by “bots” run contrary to best practices and are considered an abuse of the system. Instead, passive interventions can use data storification to crystalize a debate. If the stories so generated are both engaging and unthreatening, they can draw users to the bot’s content, thus avoiding the need for a bot to push its content into an ongoing thread. The *Excelsior!* system described here aims for unthreatening engagement by packaging its data-driven stories as comic strips that integrate two sides of a particular argument into a single visual narrative. The system’s narratives are grounded in real hashtags, which allow it to passively target its outputs at the appropriate audiences.

Keywords

storification, comic strips, irony generation, automated intervention, echo chambers, Twitterbots

1. Introduction: Why So Serious?

It has been suggested that life is a tragedy for those who feel, and a comedy for those who think. We see this dichotomy writ large on social media platforms such as Twitter, where discourse around contentious topics generates a surfeit of polarising feeling and a relative dearth of rational thought. Such platforms incentivize the articulation of short, pithy positions that prize outrage over insight, and in which interactions between opposing camps fall quickly to rancour. Nonetheless, even these rancorous exchanges may be preferable to the non-engagement with antagonistic stances that is too often observed on Twitter, for at least they expose users to multiple points of view. Instead, inward-looking, defensive structures called *echo chambers*[1] insulate disputants from interactions with those with whom they are in dispute, and contribute to the growth of factionalism and the decline of real debate on Twitter.

Bots are an oft-aligned presence on Twitter, but one benign use of Twitterbots is the generation of interventions to foster engagement between holders of opposing views [2]. Such interventions

In: R. Campos, A. Jorge, A. Jatowt, S. Bhatia, M. Litvak (eds.): *Proceedings of the Text2Story’23 Workshop, Dublin (Republic of Ireland), 2-April-2023*


✉ tony.veale@UCD.ie (T. Veale)

🌐 <https://Afflatus.UCD.ie/> (T. Veale)

🆔 0000-0003-2375-1811 (T. Veale)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

can cut to the heart of a dispute, by repackaging the nub of a conflict in an engaging form. Yet, while many follow bots out of an appreciation for their whimsical and borderline-human outputs, few users welcome unsolicited intrusions from bots in the form of direct messages, replies or explicit mentions. Indeed, even bots that politely point out spelling errors can be met with vitriol and contempt [3]. Few of us like to be lectured by strangers, least of all by bots. So, our goal with *Excelsior!* is the generation of narrative interventions that are as engaging as they are unthreatening, and which users can find for themselves via the judicious use of hashtags.

Key to this engagement is the use of comic strips as a narrative medium. As belied by the name, comics have their origins in the "funny pages" of newspapers, where they were meant to entertain more than to educate. Yet comics are a sequential art form [4, 5] that is not limited to tales of funny animals and superheroes. Here we want to exploit each of these qualities, to give narratives about serious topics a harmless comedic form that is more likely to foster engagement than suspicion and outrage. Crucially, each generated comic should balance two points of view, an argument and its converse, as articulated in the underlying data source, which in this case is the ongoing debate on Twitter about vaccines, guns, climate change and abortion.

We initially viewed each of these debates as distinct, and collected four separate corpora of tweets via Twitter's streaming API, guided by seed sets of topic-related hashtags. We realized, however, that the four debates instantiate a single master debate concerning the proper balance of power between the state and the individual, and although each corpus has unique hashtags of its own, many tags – especially those of a political nature – recur across domain boundaries. *Excelsior!* proceeds by first identifying hashtags that convey a stance toward a topic, such as *#FireFauci*, *#TrumpIsGuilty* and *#GetVaccinatedNow*, and then orders related tags into sequences of mounting emotion, such as from curiosity to skepticism to disgust. An emotional inversion is performed mid-sequence, such as from disgust to admiration, to shift the narrative to an opposing viewpoint. The complete sequence is then rendered as a comic, one panel per hashtag, that balances both points of view. This comic can then be tweeted as an animated GIF along with the hashtags that punctuate its narrative beats. This approach to narrative extraction does not seek to summarize the totality of a debate. Rather, it treats a debate as a space of viewpoints, and samples stories from this space in a way that, over time, cumulatively mirrors its emphases.

2. A Comedy For Those Who Think: Generating Comic Strips

Comics are a medium for story-telling that requires a narrative impetus. For the *Comic Chat* system of [6], this impetus comes from the interactions of online chatroom users. User texts are not summarized but placed verbatim into speech balloons above cartoon depictions of the users. Each conversational beat produces a single panel, and sentiment analysis is used to determine which variant of a user's comic avatar is associated with each speech act. But this impetus can also be machine-generated, and comics offer a viable medium for rendering automated stories, as in the story-to-comic generators of [7], [8] and [9]. This can be modeled as a text-to-text generation task if each comic is specified using XML, as in the CSDL (Comic Strip Description Language) of [7], the CBML (Comic Book Markup Language) of [10] or the ComiXML of [9]. *Excelsior!* is built on ComiXML, as it allows a comic to be specified as a specific arrangement of visual assets, drawing from a repertoire of hundreds of character poses and panel backgrounds.

This is a symbolic, componential approach to building comic strips, in contrast to the neural approaches typified by [11] and [12]. Neural approaches are trainable, and so are adaptable to specific data sets and visual genres (e.g., *Manga* in [11], *Dilbert* in [12]). They are, in principle, capable of generating diverse images to match a given text prompt, although the visual outputs of the generative adversarial networks in [11, 12] are often blurry and ill-formed. Moreover, the relationship between image and dialogue, which is the crux of the comics medium, is difficult to control in such models. This relationship is crucial when comics are used to package interventions into a debate, especially when the goal is to balance opposing points of view. Alternatively, images and text may be generated separately, by models that specialize in each. For instance, very large language models such as *GPT3* and *ChatGPT* can be used to generate stories for a given prompt [13], in the desired form (e.g., a two-person dialogue, a one-act play). To provide a suitable context to the generator, the prompt may in turn be generated by existing narrative extraction methods [14], as applied to a corpus of interest. Individual text fragments can then be used to prompt an image generator such as *Dall-E* or *Stable Diffusion* [15] to create a panel setting for each. But very large language models are resource-intensive blackboxes that are not conducive to the development of small-footprint systems; neither do they permit easy interrogation of their logical processes. A symbolic model, in contrast, ticks both of these boxes.

Excelsior! uses its lexico-semantic models to map from specific hashtags to emotion frames and topic-relative stances. These are then mapped to dialogue, poses and backgrounds for the characters that will convey them in a comic. For example, the tag *#TrumpIsGuilty* is matched to the pattern *#[person]IsGuilty* which in turn maps to the emotion *blame* and the stance *rejection* (vs. *acceptance*) toward the topic, Donald Trump. In this case, the topic is a person that can “play himself” in the comic. For a tag such as *#CovidIsReal*, the topic Covid is not a character, but is instead inserted into the dialogue of other characters. ComiXML allows characters to be configured for gender and for hair, skin and lip colour, allowing *Excelsior!* to colour-code characters by stance: pro-characters are blue-hued, while their antagonists are red. An emotion graph links emotions such as *blame* to others that can precede it or follow it in a story. Thus, *blame* can lead to *disgust*, *contempt* or *cowardice* if the data supports this escalation with a hashtag that frames the topic accordingly (e.g., *#TrumpIsaNationalDisgrace* or *#TrumpInHiding*).

Excelsior! seeks to place an emotion-driven, narrative ordering on the hashtags of the dataset. It respects the emotions inherent in the data, if not their actual ordering. The emotion graph also links emotions to their converses, allowing *Excelsior!* to perform a stance-reversal mid-comic. Thus, *blame* can prompt a reversal into *admiration*, *gratitude* and *heroism* if the data supports it (e.g. via the tags *#WeStandWithTrump* or *#TrumpsArmy*). The colour-coding of characters ensures that readers understand the reversal as a shift from one speaker (anti) to another (pro). The result is a comic-strip that balances both points of view without giving dominance to either.

3. Knowledge Representation in *Excelsior!*

As a symbolic system, *Excelsior!* crucially relies upon a range of explicit knowledge sources. These capture its understanding of comics as a visual medium and, more specifically, its understanding of social media flashpoints and hashtags. *Excelsior!*’s key resources are the following:

- **The visual lexicon:** this comprises the hundreds of characteristic poses, props and

settings that are combined (via an XML specification) to create a single comics panel.

- **The textual lexicon:** this provides a mapping of words and phrases to the unique entities which are referenced in online arguments and about which comics are generated.
- **The hashtag model:** this permits hashtags to be analyzed in terms of recurring tag patterns, such as *#Fire{personal}*. Each pattern links the referenced entity to an emotional framing (e.g., *firing*) that indicates e.g. whether the speaker is for or against the entity.
- **The dialogue model:** This suggests captions for a framing and dialogue for its participants. These include the protagonist (who offers the argument), an antagonist (who questions the argument) and any personal entities who are referenced by the argument. Captions and dialogue are modeled as reusable text templates into which argument-specific entities are inserted, such as “Wake up to the reality of *{problem}!*”
- **The visual articulation model:** These participants are assigned visual poses on the basis of what they are made to say, so that any visual metaphors in the dialogue (e.g., “Send *{personal}* packing!” or “Knock *{personal}* off their perch” are rendered appropriately.
- **The emotion model:** *Excelsior!* currently employs 94 emotion frames, from *firing* and *disgust* to *treason* and *prison*, inspired by the semantic script theory of humour (or SSTH) first proposed in [16]. These are linked by a graph that indicates which frame-shifts are allowed within the same viewpoint (e.g., *treason* → *prison*) and which indicate a radical change of viewpoint (e.g., *treason* → *heroism*). This graph allows *Excelsior!* to develop a comic as a sequence of mounting emotions, before then switching to an opposing stance.
- **The topic model:** The transition *treason* → *prison* assumes that the entity accused of treason is also the one facing jail time since, by default, *Excelsior!* only shifts between alternate framings of the *same* entity. However, a topic graph defines allowable transitions between related entities in the same comic. Thus, a comic can transition from *Democrats* to *Joe Biden*, or from *green energy* to *solar power*, without muddying its core argument.

These resources work in concert to produce a comic. Hashtags in the data are analyzed to identify the entities they reference and to understand how each is emotionally framed. These frames are sequenced by the emotion and topic models, to produce the “plot” of the comic. To render this plot, each emotion frame is mapped to a different panel of the strip using the visual articulation and dialogue models, to generate XML that calls on the assets of the visual lexicon.

4. *Excelsior!* In Action: A system demonstration

We deploy *Excelsior!* in the joint context of our four Twitter corpora, for vaccines, guns, abortion and climate change. The streaming API ensures that these corpora are constantly growing and acquiring topical new hashtags. This joint context currently comprises 2 million tweets, from which *Excelsior!* can reason about 8000 unique tags that express a stance toward a topic. The system can choose its own topic and narrative from this space, or can be directed to seek out narratives on a particular topic, such as *Covid*, *Biden* or *abortion*. The system strives for balance but not completeness in a comic. Instead, completeness is approached over time, as it repeatedly

samples the narrative space to cumulatively reflect the relative popularity of different views. A sample comic, generated for the topic *abortion*, is presented in Fig. 1. Notice how the comic balances pro- and anti-stances on abortion *across* panels, but also employs dialogue responses *within* panels to temper and sometimes challenge the assertions of individual hashtags.

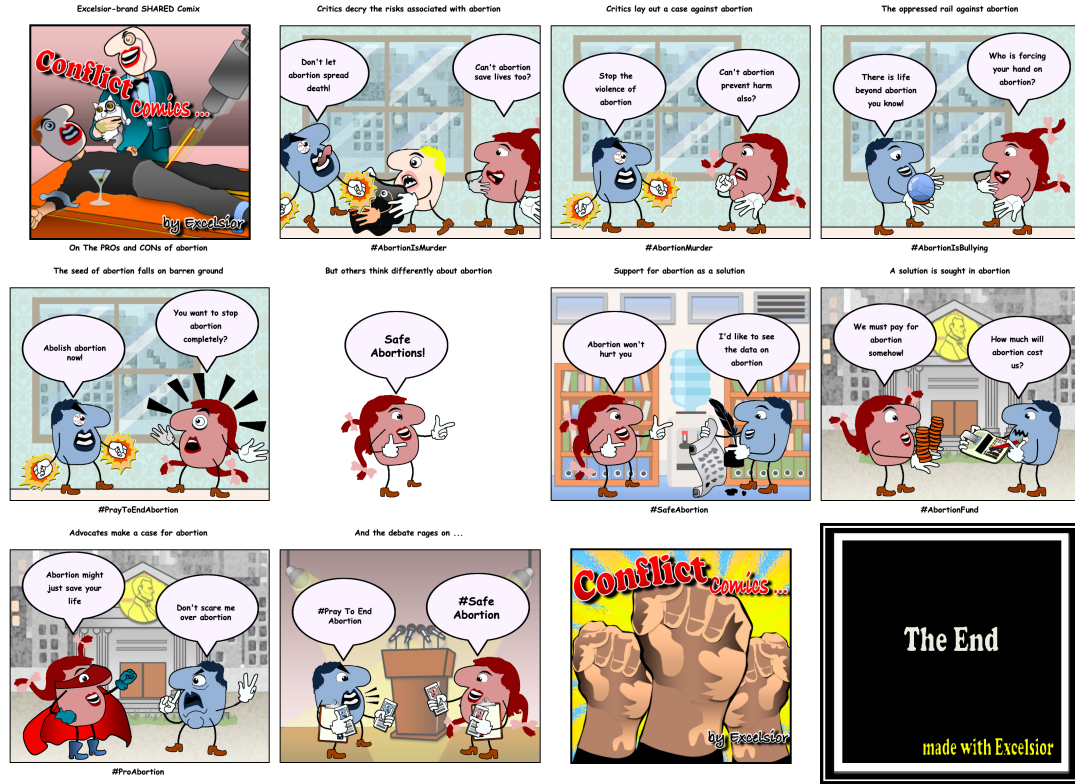


Figure 1: An *Excelsior!* comic on the topic of abortion. Stance reversal occurs in the 2nd panel of row 2.

5. Neutrality In The Balance

Since balance is key, *Excelsior!* aims for balance at three different levels of comic construction:

- **Within panels:** Each panel contains a protagonist *and* an antagonist. The former eagerly articulates a view carried by a specific hashtag, but the latter responds with skepticism. Although the protagonist drives the discussion, no view is allowed to go unquestioned.
- **Within comics:** Each strip balances the development of one viewpoint against that of its converse. Each view is developed from sequences of attested hashtags in the data.
- **Within a dataset:** The stances that are balanced within a comic may not have equal purchase in the dataset overall. The data itself may be imbalanced, so e.g., anti-vaccine voices may outnumber or out-shout the pro-vaccine alternatives. Although one side may offer fewer hashtags to work with, the two sides are always shown together in a comic.

In this balancing act, a perceived failure of neutrality at one level may be redressed at another. It may be that the hashtags conveying one viewpoint are more engaging or more trenchant than those conveying a contrary view, so that a comic appears to favour the former over the latter. It is also the case that certain visual assets are more eye-catching and humorous than others, and *Excelsior*'s choice of assets may seem to favour one argument over another. However, the tag patterns of the hashtag model, and the pose mappings of the visual articulation model, are emotion-specific but topic-independent. Those patterns and mappings are available to either side of any debate, and so other comics may well use them to convey antithetical viewpoints. For debates in which each side is equally creative, these issues will balance out in the long term.

6. Concluding Remarks

Intervention into any social media discourse, whether automated or human-moderated, is a complex business. As noted in [17], intervention strategies can greatly increase user engagement around disputed content. Yet, by drawing attention to that which we seek to curb, interventions can sometimes backfire, or provide a support structure upon which partisan users can hang further provocations. This is true of interventions that are little more than the textual equivalent of a raised eyebrow, such as "Get the facts about X" or "This claim about X is disputed," for the *ethos* of the intervener matters just as much as their *logos*. Interventions of the kind analyzed in [17] inevitably appear to support one side of the argument over another, and thus become part of the perceived power structure against which disaffected users feel they must fight. To counteract this unwanted perception of superiority and elitism, we believe that interventions must be humorous, even silly, and must be willing to mock both sides of the debate indiscriminately.

All interventions have ethical dimensions, and this is especially so of narrative interventions that treat both sides of an argument as equally valid. To be accepted as a neutral party, *Excelsior!* cannot exhibit any biases of its own, or betray those of its developers. This balance should, in principle, ensure that *Excelsior!* does not push one view over another, or cause more readers to adopt unsafe or bad-faith positions. Our aim is not to persuade, but to foster increased dialogue between opposing camps. As we explore avenues for deep-learning in *Excelsior!*, to produce more adaptive, more diverse and less predictable comics, we remain mindful of the need for the system to understand what it generates as we prepare for it to actively tweet its interventions. This will be the only test of *Excelsior*'s capabilities that really matters: do its comics make a measurable difference to how users from different camps interact with each other on Twitter?

References

- [1] P. Barberá, J. Jost, J. Nagler, J. Tucker, Tweeting from left to right: Is online political communication more than an echo chamber?, *Psychological Science* 26 (2015) 1531–1542.
- [2] C. Blaya, Cyberhate: A review and content analysis of intervention strategies, *Aggression and violent behavior* 45 (2019) 163–172.
- [3] T. Veale, M. Cook, *Twitterbots: Making machines that make meaning*, MIT Press, Cambridge, 2017.
- [4] W. Eisner, *Comics & Sequential Art*, Poorhouse Press, Tamarac, Florida, 1985.

- [5] S. McCloud, *Understanding Comics: The Invisible Art*, Harper Collins, New York, NY, 1993.
- [6] D. Kurlander, T. Skelly, D. Salesin, Comic chat, in: *Proceedings of SIGGRAPH'96, the 23rd annual conference on Computer graphics and interactive techniques*, ACM, 1996, pp. 225–236.
- [7] T. Alves, A. McMichael, A. Simões, M. Vala, A. Paiva, R. Aylett, Comics2d: Describing and creating comics from story-based applications with autonomous characters, in: *Proc. of the 20th Annual Conference on Computer Animation and Social Agents*, 2007.
- [8] R. Pérez y Pérez, N. Morales, L. Rodríguez, Illustrating a computer generated narrative, in: *Proc. of the 3rd Intern. Conf. on Comput. Creativity*, 2012, pp. 103–110.
- [9] T. Veale, Two-fisted comics generation: Comics as a medium and as a representation for creative meanings, in: *Proceedings of ICC-22, the 13th International Conference on Computational Creativity*, 2022, pp. 59–66.
- [10] J. Walsh, Comic book markup language: An introduction & rationale, *Digital humanities quart.* 6 (2012).
- [11] T. Melistas, Y. Siglidis, F. Kalogiannis, I. Manouach, A deep learning pipeline for the synthesis of graphic novels, in: *Proc. of the 12th Intern. Conference on Computational Creativity*, 2021, pp. 256–265.
- [12] B. Proven-Bessel, Z. Zhao, L. Y. Chen, ComicGAN: Text-to-comic generative adversarial network, *ArXiv abs/2109.09120* (2021).
- [13] Z. Xie, T. Cohn, J. H. Lau, Can very large pretrained language models learn storytelling with a few examples?, *ArXiv abs/2301.09790* (2023).
- [14] B. Santana, R. Campos, E. Amorim, A. Jorge, P. Silvano, S. Nunes, A survey on narrative extraction from textual data, *Artif. Intelligence Review* (2023). doi:10.1007/s10462-022-10338-7.
- [15] R. Gozalo-Brizuela, E. Garrido-Merchan, ChatGPT is not all you need. a state of the art review of large generative ai models, *ArXiv abs/2301.04655* (2023).
- [16] V. Raskin, *Semantic Mechanisms of Humor*, D. Reidel, 1984.
- [17] S. Zannettou, "i won the election!": An empirical analysis of soft moderation interventions on twitter, in: *Proc. of the 15th International AAAI Conference on Web and Social Media (ICWSM-2021)*, 2021, pp. 865–876.