

# ***THE ARTIFICIAL CARTOONIST:***

KEY CHARACTERISTICS OF AI-ASSISTED  
SEQUENTIAL STORYTELLING

ERIK BARKMAN

VIA University College, Denmark

**Corresponding author**

Erik Barkman

Email: erba@via.dk

**Erik Barkman** is an assistant professor of Graphic Storytelling at The Animation Workshop/VIA University College. He has written extensively on comics and animation since the mid-aughts for magazines like Strip!, Soundvenue, Nummer9.dk, and Stripgids. He is currently the designated comics critic of the Danish national weekly newspaper Weekendavisen, as well as a board member of the Danish Comics Council.

**Schedule for publication**

Paper Submitted: 29th September 2025

Accepted for Publication:

Published online: 30 May 2026

## Abstract

Since early 2022, AI image generators have made it possible for anyone to create comic-like sequential narratives, no drawing skills required. However, after an initial flurry of interest among creators and readers, so-called “AI comics” now seem to garner limited attention outside the realm of avantgarde experimentation.

This article offers a brief introduction to the form before moving into a close reading of two specific comics, one AI-generated, one human-made, informed by Mikkonen and Braithwaite’s concept of Figural Solidarity. Contrasting the two comics teases out characteristics of AI-generated visual storytelling which, judging by their reappearance in multiple other comics created with the help of generative AI, seem to stem from inherent limitations of the image generators themselves.

The article discusses how these characteristics may affect the reading experience and perhaps even the limited popularity of “AI comics”, and what, knowing what we now know, we may surmise about the future of the artificial or artificially enhanced cartoonist.

**Keywords:** AI, AI comics, comics analysis, sequential storytelling, iconic solidarity, figural solidarity

## Introduction

The history of AI-assisted comics production is dominated by aspiring writers using AI image generators such as Midjourney or Stable Diffusion<sup>1</sup> to illustrate comic-like sequential narratives – one variant of what Giorgio Busi Rizzi has labeled ‘algorithmic authorship’ (Rizzi, 2024). In practice, such algorithmic authorship can take a number of forms: a human-made script may be illustrated partly or wholly with images generated by an AI image generator, for instance, or both script and images can be partly or wholly AI-generated, though the latter currently tends to lead to either very avantgarde (Manouach, 2021) or nonsensical (Todasco, 2025) results. One could also imagine a sequential narrative written by an AI chatbot and illustrated by a human artist, though examples of this seem exceedingly scarce. In this article and unless otherwise specified, “AI comics” will function as a shorthand for comic-like sequential narratives produced through the currently dominant approach, wherein human-made scripts are illustrated with AI-generated images.

In the Western world of printed comics, AI comics are largely nonexistent outside of self-published books; no major publishers have put out AI-assisted comics except in the few cases where artists have been revealed to be using AI image generators in their work – typically leading to public outcry and withdrawal of the work by the publisher (Johnson, 2024). It may be the case that widely read AI comics exist, just not in print or for that matter on webcomic platforms, where they are generally met with very few views and likes. Still, their absence from search-engine results suggests that no AI comics are garnering much attention or

generating much conversation right now, online or in print, at least in the US and Europe. This may be attributed to several factors: AI comics do not have the selling points of established IP, for instance well-known superheroes, and since the creators are typically outsiders to the comics world, they don’t have any name recognition. They also tend not to have much experience writing or visualizing comics, which might negatively impact the quality of their work.<sup>2</sup> Looking beyond the Western world, however, there is some indication that AI comics might be more widespread and accepted elsewhere. In Japan, popular online manga have been revealed to be AI-generated (Zotomayor, 2025); in South Korea, the webtoon app Pocket Toons presents a dizzying array of AI-assisted comics; and in Russia, state-sponsored or aligned “Z-Comics” with AI-generated imagery are widely distributed. However, language and cultural barriers make it very hard to gauge exactly how this work is perceived in the native markets.

Moving beyond structural considerations and subjecting examples of AI comics produced in the US and Europe to a closer reading and aesthetic evaluation, it is hard to escape the suspicion that part of the reason AI comics don’t seem popular with readers could be that they are simply not as enjoyable to read as human-made comics. To pursue this supposition, the following article compares a comic produced through human-AI interaction to a similar comic created solely by humans. The analysis finds that, as dazzling as the output it produces can be on a surface level, the technology still displays severe limitations when it comes to producing compelling sequential narratives to rival the work of skilled writers and artists. Specifically, I posit that the central issue

<sup>1</sup> For an introduction to the technicalities of AI image generators and large language models (LLMs) as they relate to comics, see Rizzi, 2023

<sup>2</sup> Why established comics artists and writers dabble so little in the AI space is a whole separate question; one explanation could be that the legality of copyrighting AI art is currently in dispute, which makes it hard to monetize such work. Another likely reason is the skepticism towards the technology expressed in many corners of the comics community – as illustrated by, for example, the comments thread of Manouach, 2025 – grounded in concerns about the ethics of training LLMs on copyrighted material and the environmental impacts of the technology, an aesthetic distaste for AI-generated work, and perhaps also that artists simply don’t find prompting very satisfying compared to manual creation (Kim, Kim, Lee & Oh, 2025).

of work produced with LLMs has to do with the coherence of the visual narrative – on several levels.

I do not mean to disqualify the potential of current LLMs to produce artistically interesting work that explores the intersection of human intentionality and algorithmic creativity – in the way that the ‘synthetic comics’ of Ilan Manouach (2025) do, for instance. Rather, by analyzing AI comics that pursue the type of traditional story structures that Hollywood guru Robert McKee (1997) might call ‘arch plot’ narratives – stories that feature characters who learn and grow and which are based on an Aristotelian conception of story structure, emphasizing a clearly delineated beginning, middle, and end – and which are conveyed through visual storytelling in the tradition of the artists that have defined comics in their current form (like Hergé, Jijé, Eisner, Kirby, Tezuka, etc.), I aim to explore their ability to replace the work of professional level comics artists as feared by many creative professionals and hoped for by those dreaming of a democratization of comics production.

### AI comics: An overview

No comprehensive history of AI-generated comics as an art form – if indeed it can be considered a separate art form – exists, and it is outside the remit of this article to supply it. Suffice to say that early experiments with AI-assisted, comic-like narratives proliferated after the launch of AI image generator DALL-E in January 2021, and that the production of such narratives seemed to gain traction in earnest around Summer/Autumn 2022 – prompting overview articles on mainstream comics websites like The Comics Beat (Stern, 2023) as well as academic inquiry (for instance Rizzi, 2024).

Unlike early experiments, like Ilan Manouach’s (2021) *Fastwalkers* or Dave McKean’s *Prompt* – which were narrative experiments explicitly concerned with exploring the logic and visual ‘language’ of AI image generators and chatbots – a

number of AI comics published online around this time, like *Lungflower* by Brian Martinez (2022), *Summer Island* by Steve Coulson (2022) and *Entering the Data Core* by Christen Bach (2022), attempted to tell stories with traditional ‘arch plot’ narrative structures by adding LLM-generated images to human-written scripts.

Typically, these sequential narratives were made by creatives from outside the world of comics harboring dreams of authoring sequential narratives. For example, in the postscript to *Summer Island*, transmedia studio executive Steve Coulson describes himself as a lifelong comics lover and implies that he has always wanted to create comics of his own, but “the problem has always been – I can’t draw, which is a bit of a setback in a visual medium. And I had no friendly artistic collaborator who trusted me enough to take on a project like this (wisely)” (Coulson, 2022, p. 32). Instead, Midjourney became his ‘artistic collaborator’.

On the AI comics aggregator Alcomicbooks.com, the examples given above are categorized as “First Generation AI Comics [sic] Books”. Their style of visual narration seems largely defined by the early limitations of AI image generators. At this point in time, there were few tools available to create series of images with recurring, recognizable characters (Stern, 2023). The first-generation AI comics typically worked around this in various ways, for instance by having protagonists be hooded (Bach, 2022), unseen (Coulson, 2022), or shown only from a distance (Martinez, 2022). Further, it was even harder to dictate the composition of AI-generated images in mid-2022, making it difficult to create clear and consistent panel-to-panel progressions. As a result, these early AI comics narratives tended towards illustrated prose, using text in caption boxes to guide the readers’ understanding of images that were often semi-surreal and challenging to parse without the guidance of narration. Speech balloons were rare, and wordless passages, where narrative is conveyed entirely through images, rarer still.

The “Second Generation AI Comics [sic] Books” on Alcomicbooks.com, as described on the website, are characterized by “using advanced AI technologies”. What this means more specifically is not elucidated, but most likely refers to then-new methods for inserting recurring characters into images, dictating poses and compositions by, for instance, uploading sketches to the LLM, and more – all making possible the creation of sequential narratives that are less reliant on prose text to explain the events of the story. The ‘second generation’ AI comics also showcased the increasing sophistication of image generators when it came to rendering images in different, albeit generalized artistic styles – bringing to mind superhero comics or manga, for instance.

Newer works on Alcomicbooks.com tend to emulate US style comic books featuring DC/Marvel style superheroes or Vertigo/Image Comics style genre narratives: horror, science fiction, fantasy. At first glance, they bear a convincing similarity to the kind of professionally produced work found on the shelves of American comics book shops, but they seem to enjoy nowhere near the visibility or popularity. AI comics are hard to find online – at the time of writing, a simple Google search yields a plethora of “AI comics generators” before Alcomicbooks.com finally appears on page two of the search results. Sorting the books on that site by ‘most popular’ brings R. Neos *Sojourn Realm* to the top of the list. But while Alcomicbooks does not allow users to like or comment, an Amazon listing of the print on demand version of *Sojourn Realm* shows zero ratings or comments. Similarly, the newest e-book on Alcomicbooks at the time of writing – and the fifth ‘most popular’ – is *Tokyo Roommates 3*. It can also be found on online comics platform Globalcomix.com, where it currently has only 231 views.

## Methodology

In preparation for writing this article, I studied a number of AI comics via social media links, overview articles,

and aggregators like Alcomicbooks.com and the recently launched Storinex.com. I found striking similarities in the type of visual storytelling employed in these comics and ultimately selected one that I found representative of the form. In the following, we will compare it to a human-made comic to highlight key aspects of the AI comic by contrast. Then, we will briefly survey several other AI comics to get a sense of the extent to which our findings are generalizable.

The AI comic in question is chapter 5 of *Classic Epic* (Allen, 2025a), adapted from the 16<sup>th</sup> century Chinese epic *Journey to the West* – commonly attributed to Wu Cheng'en – by graphic designer Ryan Allen using the image generator Stable Diffusion. Allen wrote the script based on the original text and then prompted panels individually before organizing them into pages. *Classic Epic* was selected because it is relatively current (January – February 2025) and so made with fairly up to date technology, and because Allen (2025b) has shared the process behind the comic, removing some of the need for guesswork that can come with analyzing AI comics – specifically: how have the images been prompted and in what ways and to what degree have they been digitally altered post-generation.

The (for lack of a better term) human-made comic is issue 5 of the 10-issue miniseries *The Moon is Following Us* (Johnson & Rossmo, 2025) by Daniel Warren Johnson, credited as ‘Creator, Writer, Artist’, and Riley Rossmo, ‘Creator, Artist’. Coloring is credited to Mike Spicer and lettering to Shawn Lee. *The Moon is Following Us* is not a mythological epic like *Classic Epic*, but it is a fantastical tale taking place in an imaginary space – specifically the dreamscape of a young girl. It was published at the same time as *Classic Epic* chapter 5 (February 2025), and like that AI comic, it is created digitally in a semi-realistic US comic book-style visual idiom.

As the theoretical framework of the analysis and especially the discussion of the impact of visual storytelling on reading

experience that will follow, I am taking as my starting point *Figural Solidarity: Grappling with Meaning in Comics* by Jean Braithwaite and Kai Mikkonen (2022). Working from Thierry Groensteen's (1999) concept of iconic solidarity as presented in his influential *Système de la bande dessinée* and drawing on their own findings in eye tracking studies – where readers were shown to fixate on recurring characters, following them around from panel to panel – Braithwaite and Mikkonen propose focusing on reappearing figures rather than panels as the key to making sense of (traditional) comics narratives<sup>3</sup>, coining the term “Figural Solidarity”. “A figure” in Braithwaite and Mikkonen's definition is “any discernible diegetic entity that the comics readers can pick out from among all the marks inside the comics panels” (para. 28). Further, “salient figures” are “just those figures parsed by the reader as significant for narrative processing purposes” (para. 47) – so not random background characters, but entities relevant to the progression of the story being told. Finally, they argue that “salient figures” are singled out by the narrator and recognized by the reader primarily by virtue of being “repeated/recognizable”.

## Analysis

In chapter 5 of *Classic Epic*, the Monkey King, seeking to become immortal, has begun studying under the immortal bodhisattva Master Subhodi. The first three pages are taken up by a montage of activities performed by our main protagonist while under the tutelage of the Master, for instance on the second comics page of the chapter: “He discussed scriptures and doctrines with [his schoolmates]. He exercised, practiced painting, and performed various meditations.” Looking at this page, what is most immediately striking is how the panels are all composed in the same way: The Monkey King is in the center of the image, often taking up most of the visual real estate.

He is always drawn at eye level, there are no bird's or worm's eye views, and is shown in either full shot (full body visible), medium shot (upper half of the body) or medium close-up (chest and/or shoulders plus head)<sup>4</sup> – never closer and never from farther away. His face is always in the panel, there are no images showing only his legs or his torso and hands, and he is always facing the reader. As we will see later, this is generally an overrepresented type of image composition and cropping in AI comics generally, which would seem to suggest that this approach to constructing images is favored by LLMs.

Looking closer at the individual panels, three out of five show curious mismatches between the actions described and the images shown. The panel purportedly showing the Monkey King exercising simply sees him holding a rope while smiling peacefully. “In his leisure time, he would sweep the grounds” is accompanied by a wide panel of his face and the suggestion of a broom to his left – surely a close up of the actual sweeping would have been more illustrative? – and “He discussed scriptures and doctrines with [his schoolmates]” shows him writing in what looks like an already written book, a face in the background standing in for his schoolmates. No discussion seems to be going on. Similarly, as suggested in the description of the exercise panel, facial expressions often seem mismatched to the events in the panels with the exception of the one in which the Monkey King meditates – mostly, they default towards neutrality – and he is looking out at the reader in three of the four panels, even as he is supposed to be focusing on exercising, painting, and sweeping. Conceivably, these panels reflect the still-significant challenges of writing prompts that produce the exact images desired, especially when it comes to facial expressions conveying nuanced and varied emotions (Pagan, 2025).

<sup>3</sup> “To us it seems clear that a semantics of comics has to be willing to formalize units below the panel level. So, even without a solid foundation beneath us, we will simply have to start somewhere in the middle, wherever we can get a toehold. This is the value of figures” (Braithwaite & Mikkonen, 2022, paragraph 41).

<sup>4</sup> For lack of a specialized vocabulary for the *mise-en-scène* of comics, I am here using terms borrowed from filmmaking.



Figure 1 *Classic Epic*, page 2 of chapter 5 (Allen, 2025)

Finally, it is striking how differently stylized the Monkey King is from panel to panel. Compare, for instance, the lush rendering of the panel in which he is painting – multiple gray tones applied in a manner reminiscent of water color or gouache, fur carefully delineated and textured – to the much simpler style of the neighboring meditation panel, reminiscent of traditionally inked comics panels: simple outlines filled in with flat colors, shading in a single tone, next to no texture applied to the fur. Simply put, these panels look like they're from different comics, drawn and colored by different artists. The inconsistencies extend to the design of the Monkey King himself: in the exercise panel, his face is much wider than in the other panels, and the fur on his chin is sometimes goat-tee-shaped, sometimes almost entirely rounded<sup>5</sup>. Similarly, the backgrounds are different in every panel; though they do display some of the same elements (mountains, a lake), no stable sense of setting is established.

Reading through the rest of chapter 5, many further examples of stylistic inconsistency can be observed: the immortal bodhisattva under whom the Monkey King studies is depicted in a painterly style devoid of outlines as if he walked in from a religious painting, and when he gives a lecture, the accompanying montage looks as if each panel was drawn by a different artist, ranging from hyperdetailed and painstakingly rendered line drawings to Mike Mignola-ish figures with blocky black-spotting – and from semi-realism to manga-like stylization. This goes for inconsistency of design, too. In some panels, the Monkey King looks like a skinny, fresh-faced youngster; in others, he is burly, weathered and ripped. And as for the environments, their arbitrariness is particularly noticeable in chapter 4 of *Classic Epic*, when the Monkey King arrives at Master Subhodi's temple; he reaches it through a door in a mountain, but the door is never placed in the context of a larger environment, and there is no logical spatial transition from

the mountainous region in which the door is found to the lush temple grounds on the other side of the door.

Organizing our observations about the visual storytelling of *Classic Epic* into broad categories, we find that it is characterized by a *mise-en-scène* which is limited to a fairly narrow range of zoom levels, angles, and compositions; character posing and expressions that are noticeably unspecific and imprecise relative to what the prose suggests should be conveyed; and stylistic inconsistency in terms of drawing and rendering style as well as character and environment design. These characteristics become especially noticeable when comparing *Classic Epic* to a human-made comic.

The fifth issue of *The Moon is Following Us* opens with a sequence focusing on the character Brio. The premise of the story is that the girl Penny has fallen into a deep sleep from which she cannot seem to awaken, and her parents discover that she is being held captive by her dreams. They find a way to enter her dreamscape and discover allies there, including a dream version of Penny's childhood plush frog Brio. The first page of the issue shows Brio lying in a crosswalk (seen directly from above) – this is how Penny lost her toy as a toddler: she dropped him while crossing the street – and then moves to a montage of an action hero version of Brio readying himself for battle against the forces holding Penny hostage in her dreams. Then, we move to a closeup of the dropped plushie before ending the page on a long shot of Brio standing in front of Penny's parents. Comparing this to the *Classic Epic* page, we see a much wider variety of zoom levels, including several showing Brio's hands or torso and extreme long shots showing the plush frog. We also see a greater variety of angles, from the extreme bird's eye view of the crosswalk to the subtler worm's eye view of the final panel. And even when compositions are centered, the salient figures are pushed slightly


<sup>5</sup> Incidentally, in his FAQ (Allen, 2025b), creator Ryan Allen reports having mitigated worse facial hair inconsistencies by hand, digitally painting on top of AI-generated images in several instances.



THINGS CHANGE  
WHEN YOU GET RUN  
OVER BY A CAR.



DROPPED ME IN THE MIDDLE  
OF THE CROSSWALK.



I KNOW  
IT WASN'T  
PENNY'S  
FAULT.



SHE  
WAS JUST  
A BABY  
THEN.



BUT THE TRUTH IS  
THE TRUTH. NOTHING  
IS FOR CERTAIN.



LOOK WHAT  
HAPPENED TO  
OUR BEAUTIFUL  
QUEEN ASHLEY.

WHO KNOWS  
WHAT'S IN STORE  
FOR US NEXT?

Figure 2 The Moon is Following Us, page 1 of issue 5 (Johnson, D. W., & Rossmo, R., 2025)

to one side, as in the top image, where plush Brio is to the left of the center, balanced out by the car on the right.

As for the physical action, Brio is shown performing easily decipherable and precisely delineated tasks – loading ammunition into a gun, grabbing grenades, holding his crown – with no need for captions to explain what is going on. And his ‘acting’, while not conveying much in the way of emotion beyond the grim determination of dreamscape Brio contrasted with the naïve smile of plushie Brio, still does this in a clear and consistent manner. Finally, there’s a uniform stylization on the panels set in the present day of the narrative, chunky linework supplemented with a textured, loose rendering of light and shadow. In the two panels showing Brio’s ‘origin story’, rendering is done with inky hatching, and the colors are sepia toned, but this stylistic break has an explicit storytelling purpose: to show events set in a different time and place. Similarly, the designs of recurring characters (past and present Brio) and environments (the sidewalk) are consistent throughout.

Compared to *Classic Epic*, the *mise-en-scène* of *The Moon is Following Us* shows significantly greater variation in zoom, angles, and compositions; character posing and expressions are more specific and reflective of what is going on in the narrative; and the drawing and rendering style as well as character and environment design are not comparably stylistically inconsistent. In other words, our observations about the visual storytelling of Allen’s AI comic seem to be confirmed by the comparison.

Of course, *Classic Epic* is only one AI-assisted comic out of many, and it is entirely reasonable to question to what extent other examples of the form display the same limitations or, in more neutral terms, characteristics in terms of visual storytelling. Todasco (2025), makes for an interesting comparison, because in contrast to Ryan Allen’s work on *Classic Epic*, Todasco has prompted the ChatGPT-4o image generator – best known for viral Studio Ghiblifications of family photos

and profile pictures – to generate an entire “Silver Age/1960s” superhero comic, from script to layout to finished pages, based on a single prompt. As such, the results could be argued to show the image generator’s go-to solutions for visual storytelling in an unfiltered way, unassisted by a human artist in terms of granular prompting of individual panels, assembly of layouted pages, and post-generation edits.

Looking at the four pages of *Captain Coldwar vs. The Nuker*, we find the same tendency regarding *mise-en-scène* towards centering characters and using a narrow range of zoom levels and angles as in *Classic Epic*, though there is a suggestion of a top-down view on page 3, panel 1. There are also more panels featuring multiple characters, and these by necessity don’t center the main protagonist – instead, the ‘salient figures’ are generally distributed so that the distance between them and the edge of the frame is equal, i.e. somewhat mechanically. Character posing and expressions are hard to judge, since the story itself is close to nonsensical and not much is required in terms of ‘acting’ – a limited number of facial expressions are used repeatedly, though to be fair, that is not entirely unlike what one might see in a typical superhero comic from the 1960s. The stylistic inconsistency of Todasco/ChatGPT-4o’s comic is markedly less than that of *Classic Epic* in terms of rendering, but much more pronounced when it comes to character design: As remarked by Todasco (2025) in his own evaluation, “The protagonist seemed to physically morph on each page, and his identity shifted from Captain Coldwar to Super-Patriot to Cosmic Crusader” (para. 10). *Captain Coldwar vs. The Nuker* is interesting in that it could give some indication of how the typical image generator’s approach to making images might be affecting the visual storytelling of AI comics with much more significant human involvement. However, it should be noted that the “Silver Age/1960s” superhero comics Todasco asked ChatGPT-4o Image Generation to imitate do tend more toward the theater-like *mise-en-scène* seen in *Classic Epic* than the more cinematic visual storytelling of *The Moon is Following Us*.

Turning to examples of comics written by humans and illustrated by LLMs, we might look at the top three “most liked” titles on Storinex.com, which at the time of writing seems to be the most active platform for English-language AI comics. *Galaxy Girl* (Wizard\_of\_Ozymandias6, 2025) tops the list, a space opera centered around the exploits of a young, buxom heroine. The newest chapter at the time of writing, *Aquarius*, showcases a slightly more flexible approach to *mise-en-scène* than in the other examples: Characters are not always centered and are occasionally shown from above, though only when they are lying down. There are also a few more zoomed-out images and some closeups, the latter most likely achieved by cropping generated images. As for character expressions, they remain noticeably imprecise – on page 12, for instance, Galaxy Girl is undergoing a stress relief treatment that seems to be bringing her to a sort of climax, and yet her face remains stuck in a doll-like expression of slight bafflement. Stylistic inconsistency is a factor, too, for instance on page 10, where Galaxy Girl’s hair is rendered in two very different ways in the first two panels. That page as a whole is also a good example of how the rendering shifts between an “inked”, outline-based style and a digitally painted style, seemingly to no narrative end – much like in *Classic Epic*. Next on the “most liked” list is *Synaptik* (Express\_00, 2025). The series’ creator experiments with using silhouettes in his panels, but other than that, the same lack of variation in the *mise-en-scène* as regards zoom, angles, and compositions as we have seen in other examples is on full display in the currently newest, sixth chapter; the characters are inexpressive; and while character designs seem fairly consistent, the rendering style once again swerves between flat colors and realistic light and shadow. Number three “most liked” AI comic, *Red Core Resurgence* (Nomednomel, 2025), exhibits the same characteristics – some more clearly than others – and so it goes on down the list.

## Discussion

So far, we have analyzed an example of a sequential narrative illustrated with AI-generated images and, comparing it to one visualized by human artists, teased out key characteristics of AI-assisted visual storytelling. We then confirmed our findings by looking at other examples of AI comics. In the process, the characteristics of AI comics have typically been described in somewhat negative terms – “limited”, “narrow” – working from the implicit assumption that a more complex approach to visual storytelling makes for a richer text. But of course, it is not a given that just because a seasoned reader of comics might find the visual storytelling of AI comics flawed or limited in different ways, the average reader will come to the same conclusion. In other words, even if we may accept the assertion that AI comics are less advanced or sophisticated in terms of visual storytelling and exhibit certain telltale quirks, we cannot necessarily conclude that this might hurt their readability. After all, plenty of popular visual narratives are very simply told – from *Peanuts* to *Dog Man*. In considering how the visual storytelling of AI-generated comics may affect reading experience, Mikkonen and Braithwaite’s conception of Figural Solidarity comes in handy, not least because it’s derived from eye tracking studies of the actual comics reading experience.

Applying the concept to our selected pages from the two comics, we see that both clearly present us with recurring characters for the reader to follow as they parse the narrative. In the *Classic Epic* page studied, the Monkey King is front and center in every panel, impossible to miss; in *The Moon is Following Us*, Brio similarly recurs in every panel, but the reader is put to work identifying him – in some panels we see him as a real life plush frog, in some as a cigar-chomping dream version of that plushie; some show his full body, some just his hands, a third of his face, his head and shoulders in silhouette, etc. We could

6 Many creators of AI comics work anonymously under pseudonyms like this – in line with general practice in the field of “popular AI art” as described in Philippsen, 2025

say that the metonymic leaps required of the reader are significantly more challenging in the 'human-made' comic, in which we are asked to mentally connect a past version of Brio to his current dream iteration and images of various body parts to those of his entire body. By contrast, *Classic Epic* only requires that we can recognize the Monkey King's face, shown in full and from the front on every single panel.

One could argue that *Classic Epic* and AI comics like it should be easier to read than works like *The Moon is Following Us*, requiring less of the reader thanks to their limited variation in zoom level, angles, and image compositions. On the other hand, it may be that comics readers enjoy a challenge. As Scott McCloud (1993) argued in *Understanding Comics*, sequential narratives come alive in readers' minds when they imagine what happens in the gutters between panels – i.e. connect the dots of events depicted on the page in their own imaginations. This, he claimed, is how comics invite "participation" from readers, fostering "an intimacy surpassed only by the written word" (p. 69). McCloud's concept of 'closure' has proved both inspirational and controversial in comics scholarship (Oppolzer, 2023), but it seems reasonable to assume that metonymic leaps like those made in parsing the page from *The Moon is Following Us* activate the reader's imagination, allowing them to make connections between images in their minds – from parts of Brio to the totality of Brio and from past Brio to present Brio. Expressed in the terms of Mikkonen and Braithwaite: while "repeated/recognizable" characters are essential to the narrative processing of comics, it may be that readers become more engaged in visual narratives when challenged to recognize those characters across different modes of depiction.

It is debatable to which extent the narrower range of variation in *mise-en-scène* and the lack of precision in the delineation of character expressions and actions observed in AI comics has to do with the limitations of AI image generators – and to what extent it may be ascribed to a lack of ability or experience with sequential storytelling among creators of AI comics. After all, the cartoonists behind a series like *The Moon is Following Us* have years of experience telling stories in comics form between them, whereas someone like *Classic Epic* creator Ryan Allen is a relative newcomer<sup>7</sup>. If an experienced cartoonist who was also adept at prompting image generators were to create an AI comic with a straightforward narrative, the results could go some way towards answering this question. Unfortunately, I am not aware of any current examples of this<sup>8</sup>, though I do know of at least one study pitting the visual storytelling of LLMs against that of untrained human artists and awarding the victory to the human amateurs (Önal, Sariel, Sezgin & Akleman, 2025).

However, an aspect of *Classic Epic* which can almost certainly be ascribed to the way image generators function is the stylistic inconsistency – or variation, if you will – evidenced in the comic. Of course, stylistic variations can "form a structuring element" of comics and thus be "essential to communicating [their] meaning" (Brookes, 2024, p. 1); in *The Moon is Following Us*, different rendering styles are being used to show past and present versions of Brio the plush frog. But when this happens, the variation will be consistent: other flashbacks in the series are drawn in the same way. There is no similar narrative justification for the stylistic variation in *Classic Epic* and other AI comics. Less trained visual artists will often struggle with keeping characters 'on

7 As Django Beatty puts it in Beatty, 2025: "When you know what good looks like, you can see exactly how far the AI is from where it needs to be. Expertise isn't just knowing tools - it's years of building a mental map of quality. A designer sees why the typography fails. A developer sees why the architecture won't scale. A writer sees why the voice is wrong."

8 "Current", because veteran webcomic artist Ursula Vernon's *A Different Aftermath*, the only example of which I am aware, was made in September 2022 with a 'first generation' image generator.

model<sup>9</sup> and stylizing different pictorial elements so that they feel like they belong in the same visual universe: cars may be less confidently rendered than humans, for instance. And stylistic variation without a clear narrative purpose is certainly not unheard of in the 'professional' world of comics: monthly superhero comics might rotate artists, making for noticeable stylistic variation from issue to issue, manga assistants may draw backgrounds or secondary characters in a style distinct from that of the main author, and in the classic British anthology *2000AD*, each chapter of a narrative will in some instances be drawn by a different artist. Even so, I would argue that the type of panel-to-panel stylistic variation seen in *Classic Epic* is not something comics readers are currently acclimatized to.

Usually, when reading a comic produced by human hands, it is possible to get some sense of the aesthetic preferences, the strengths and limitations, and (for lack of a better phrase) the artistic temperament of the artist(s) behind the images. When reading LLM-generated comics, this sense is very much absent, as of course it must be: there are no human hands behind these images. Instead, they are synthesized from a vast swath of pictorial representations of monkeys and bodhisattvas, from cave paintings to photos, and presented in an approximately US comics-like idiom extracted from thousands of sequential narratives emerging from that tradition. It seems likely that AI image generators will develop towards delivering more stylistically consistent images over time, but it seems just as likely that it is inherently in the software's nature to generate imagery which represents the throughline of a lot of different artists' output and is thus devoid of the quirks and idiosyncrasies by which the work of the individual human artist may typically be recognized. Whatever the future may hold, stylistic inconsistency is a core trait of AI comics as things stand, noticeable to varying degrees in all the work we have discussed.

It may be that the stylistic inconsistency of AI comics, here specifically *Classic Epic*, is another explanation for their limited popularity here in the West. Returning once more to the concept of Figural Solidarity and 'repeated/recognizable' figures, it's not that readers will be genuinely confused about whether the Monkey Kings in the panels of the page examined above are the same character – if nothing else, the caption text makes it clear that he is – or if they belong in the same comic. But even if the character is identifiably the same across the different images, it is still noticeably unstable in its representation. As such, the depictions of the Monkey King could be said to become subtly uncanny or *Unheimlich* in the Freudian sense (Freud, 1919): They are almost the same, almost familiar from before, but not quite – just a little bit off. Another way of framing this phenomenon would be to posit that the unstable representation of the characters and world of *Classic Epic* creates an unintentional form of Brechtian *Verfremdungseffekt*, arguably unsuspending the readers' disbelief and alienating them from the narrative.

It is worth considering whether the concept of Figural Solidarity might be meaningfully supplemented with one of 'stylistic solidarity'. If so, this would be a concept born of AI comics: in a pre-AI world, where individual issues or chapters of comics were almost always created by the same artist or teams of artists, stylistic consistency on a panel-to-panel and page-by-page level was almost a given, rendering its importance to the reading experience mostly invisible. It may be that just like readers depend on recognizable and recurring characters to follow a (traditional) comics narrative, they depend on stable representations of characters and environments in terms of both visualization and design – i.e. stylistic solidarity – to fully immerse themselves in that narrative. If so, *Classic Epic* and other AI comics might unintentionally show the way towards new avenues in comics analysis and theory, focusing on the pictorial universes represented in comics and how

9 "Drawing characters on model" is an animation term meaning drawing characters so they look the same from frame to frame, pose to pose – or in this case, panel to panel and page to page.

their design, stylistic representation and internal consistency – their visual worldbuilding – inform readers' experience of and engagement with sequential narratives.

Applying the concept of Figural Solidarity to our observations of the differences between the two comics examined allows us to suggest two (of possibly several) reasons why AI comics do not seem to reach a substantial readership: That they are, in a sense, too easy to read and thus the act of following salient figures across the page does not allow the reader to connect the dots of the narrative in their mind and thus engage themselves deeply in the narrative; and that the lack of stylistic solidarity of sequential images produced by LLMs make them subtly *Unheimlich* and distancing, prompting an unsuspension of disbelief in readers. Adding to this – but not related to Figural Solidarity *per se* – is the lack of precision in the depiction of physical actions and emotions, which make the events of the comics harder to interpret and arguably also renders the characters more difficult to relate to and empathize with<sup>10</sup>.

## Conclusion

Compared to human-made comics, sequential narratives with AI-generated visuals seem to tend toward a less dynamic and varied *mise-en-scène*, a less precise conveyance of narrative through character posing and expressions, and a stylistic inconsistency that renders representations of the world and characters of the story unstable and arguably causes an unintentional *Verfremdungseffekt* in the reader – possibly explaining the lack of popularity of AI comics as a whole. It is debatable to what extent the visual storytelling of *Classic Epic* and similar AI comics are defined by the capabilities of LLMs and to what extent it's the creators' abilities that render them arguably less readable than human-made alternatives; further, it is hard to tell at this point in time to what degree the

stylistic inconsistencies of a comic like *Classic Epic* could be eliminated with the help of future, more powerful AI image generators. But given the current limitations of the technology, image generators as 'artistic collaborators' (Coulson, 2022) for writers and would-be writers looking to create their own sequential narratives seem to produce less sought-after comics than accomplished human artists. In the event that this does not change over time, it may be that – as Django Beatty (2025), CEO of technical consulting firm Fluxus, recently posited – "The future isn't humans being replaced by AI" (para. 70). Instead, he asserted, AI functions as a "cognitive exoskeleton", that "amplifies your strength but requires your constant guidance" (Beatty, 2025, para. 50).

Applying the terminology of Eric Kluitenberg and the archaeology of imaginary media (Kluitenberg, 2011) to the notion of AI image generators as 'artificial cartoonists', we could theorize that this is an example of an "impossible machine" (Kluitenberg, 2011, p. 66) – a construct reflective of the dreams and desires of those using or wanting to use the software as a way to finally realize their dreams of becoming visual artists without having to enlist the help of human creatives or acquiring the necessary skills themselves. There seems to be a good chance that we will look back on this conception of generative AI as being able to replace human artistic labor entirely as an 'imaginary medium' – similar to how, to use one of Kluitenberg's many examples, the early Internet promised to be an avenue for completely egalitarian communication between humans of all stripes; but in reality, "the persistence of social and cultural divides and the intricate difficulties of communication with the other have turned out to be far more tenacious than originally conceived in the early stages of the development of Internet as a public medium" (Kluitenberg, 2011, p. 61).

Conceiving instead of genAI as a 'cognitive exoskeleton' in the context of comics production, we can imagine different

<sup>10</sup> As asserted in Önal, Sariel, Sezgin & Akleman, 2025, "Effective storytelling relies on creating appropriate human affects, such as emotions, mood, or attachment".

benefits accruing from the technology, like software harnessing the power of AI to assist in more mechanical aspects of comics production. The amount of tasks that can be meaningfully automated in comics are arguably limited, since almost every word, line and color choice reflect artistic decision-making, but some promising attempts have been made, for instance when it comes to 'flattening' and adding shadows as part of the coloring process<sup>11</sup>. While not yet truly useful, these types of applications could potentially lead to more effective production and quality of life improvements for artists, and it may be that this is where AI and comics can and will intersect most meaningfully – not in the wholesale replacement of artists with experience in and a deep understanding of how to tell stories with pictures, but by freeing them from routine or less creative tasks and allowing them to produce even more accomplished and compelling narratives.

## References

- Allen, R. (2025a). *Classic Epic, Chapter 5*. Classic Epic. <https://www.classicepic.com/p/teacher>
- Allen, R. (2025b). *Classic Epic, FAQ*. Classic Epic. <https://www.classicepic.com/p/faq-what-are-the-steps-to-make-this>
- Bach, C. (2022). *Entering the Data Core*. Rock & Roll Animation. <https://christenbach.de/portfolio/entering-the-data-core/>
- Beatty, D. (2025). *Why Your AI Never Works on the First Try*. Fluxus. <https://fluxus.io/article/why-your-ai-never-works-on-the-first-try>
- Braithwaite, J., & Mikkonen, K. (2022). Figural Solidarity: Grappling with Meaning in Comics. *Comicalités*. DOI: <https://doi.org/10.4000/comicalites.7729>
- Brookes, G. (2024). Stylistic co-existence and the chronotope in *Stone Fruit* and *Onwards Towards our Noble Deaths*. *Journal of Graphic Novels and Comics, Vol 15/3*, 379–398. DOI: <https://doi.org/10.1080/21504857.2023.2262001>
- Coulson, S. (2022). *Summer Island*. Campfire. <https://www.campfirenyc.com/comics/>
- Express\_00 (2025). *Synaptik*. Storinex. <https://storinex.com/series/693afc62-c542-4852-914f-b95267148ce1>
- Freud, S. (1919). Das Unheimliche. *Imago, Zeitschrift für Anwendung der Psychoanalyse auf die Geisteswissenschaften V*, 297–324.
- Groensteen, T. (1999). *Système de la bande dessinée*. PUF.
- Johnson, D. W., & Rossmo, R. (2025). *The Moon is Following Us #5*. Image Comics.
- Johnston, R. (2024). *Accusations Over DC Comics Artists Using Artificial Intelligence Mount*. Bleeding Cool. <https://bleeding-cool.com/comics/accusations-over-dc-comics-artists-using-artificial-intelligence-mount-ai/>
- Kim, Y., Kim, S., Kim, Y., Lee, S., Oh, U. (2025). "I'm not thinking anymore, just following the path.": Investigating Task Delegation Trend of Author-AI Co-Creation with Generative Ais. arXiv preprint arXiv: 2502.16740v1
- Kluitenberg, E. (2011). On the Archaeology of Imaginary Media. In E. Huhtamo (Ed.). *Media Archaeology: Approaches, Applications, and Implications* (pp. 48–69). University of California Press.
- Manouach, I. (2021). *Fastwalkers*. La Cinquième Couche.
- Manouach, I. (2025). *In Defense of Synthetic Comics*. The Comics Journal. <https://www.tcj.com/in-defense-of-synthetic-comics/>
- Martinez, B. (2022). *Lungflower*. Globalcomix. <https://globalcomix.com/c/lungflower>
- McCloud, S. (1993). *Understanding Comics*. HarperPerennial.

<sup>11</sup> Both examples of avenues currently being explored, though so far not entirely successfully.

McKee, R. (1997). *Story: Substance, Structure, Style and the Principles of Screenwriting*. HarperCollins.

Nomednomel (2025). *Red Core Resurgence*. Storinex. <https://storinex.com/series/80edb21d-a72c-49f6-b884-8851976261d1>

Önal, U., Sariel, S., Sezgin, M., Akleman, E. (2025). Comparing Human and AI Performance in Visual Storytelling through Creation of Comic Strips: A Case Study. arXiv preprint arXiv:2507.18641

Oppolzer, M. (2023). Coherence in Comics. An Interdisciplinary Approach. *Closure. Kieler e-Journal für Comicforschung* #9.5, 7–18. <http://www.closure.uni-kiel.de/closure9.5/about>

Pagan, A., Loveland, K., Acierno, R. (2025). Evaluating the emotional accuracy of AI-generated facial expressions in neurotypical individuals. *Discover Computing Vol 28*, article 116.

Philipsen, L. (2025). Aesthetic Protocols of Popular AI Art. In T. Jan-Noël & W. Lukas R. A. (Eds.). *AI Aesthetics - AI-Generated Images between Artistics and Aisthetics*. Routledge.

Rizzi, G. B. (2023). Do AIs dream of electric comics? Generative AI models, digital memory, and creativity. In Carbé, E., Lo Piccolo, G., Valenti, A., Stella, F. (Eds.), *La memoria digitale: forme del testo e organizzazione della conoscenza: atti del XII convegno annuale AIUCD* (pp. 237–244). Università degli Studi di Siena.

Rizzi, G. B. (2024). The Networked Author: looking at contemporary authorship through postdigital comics. *ANGLICA* 33/2, 115–141.

Stern, C. (2023). *Comics with AI art: what's out there, what's to come*. The Comics Beat. <https://www.comicsbeat.com/all-about-ai-comics/>

Todasco, M. (2025). *Can AI Create a Comic Book? Volume 3: How ChatGPT-4o Image Generation Will Transform the Creative*

*Process*. Medium. <https://medium.com/@todasco/can-ai-create-a-comic-book-99d12e5a03c2>

Wizard\_of\_Ozymandias (2025). *Galaxy Girl*. Storinex. <https://storinex.com/series/280ff45e-51110-4830-be31-eb246653b42d>

Zotomayor, C. (2025). *AI-generated manga becomes top-ranked in Japan's biggest e-book store*. Automaton. <https://automaton-media.com/en/news/ai-generated-manga-becomes-top-ranked-in-japans-biggest-e-book-store/>