Fabula: An Interview with Piotr Mirowski

Piotr Mirowski
piotr.mirowski@computer.org

Alice Barale
alice.barale@unimi.it

Can AI help tell stories? And if so, how can it do so? Is it better for the "writing assistant" to be fully controllable and transparent, or is it more useful for it to be able to surprise us? These are some of the questions that emerge from this interview with Piotr Mirowski on Fabula, an AI tool that enables the creation and editing of stories through large language models.

Fabula: An Interview with Piotr Mirowski

Piotr Mirowski
piotr.mirowski@computer.org

Alice Barale
alice.barale@unimi.it

This interview with Piotr Mirowski focuses on Fabula, an AI tool l that enables the creation and editing of stories through large language models. Fabula was developed by Piotr Mirowski in collaboration with Richard Evans and a team at Google DeepMind and is currently in an experimental phase[292].

Alice: Let's start with a question that many people ask. There is a widespread fear that AI might replace human beings. Do you see Fabula as a tool that could replace human storytellers?

Piotr: Thank you for this question, which is very much on my mind. When building AI tools, there is always a tension between a logic of replacement and a logic of augmentation or support. And I am definitely in the camp of trying to support and help and potentially augment the existing human process if needed, as opposed to replace anyone or anything. So the Fabula tool was built in the logic of providing support for writers of various skill levels. For writers who are more novice, or let's say amateur, the tool can easily help generate a structure or even ideas of a script, with the objective of showing what is possible and to allow the writer to work on a script via convergent iteration. Convergent iteration is this idea that you don't necessarily start with a blank page, but rather you have already an option that is presented to you based on your original idea, and then you gradually sculpt it or shape it, according to what you want to achieve. For someone who already has existing material or for more confident writers, the tool could also provide with feedback on their existing text. For instance, you can upload your script into Fabula and see an analysis of its structure, and then explore the "what if". What would happen if you changed this character? Or if you added a new scene? And if you changed some elements of a narrative? How would everything else fall into place? You can use the tool to get a fresh look at your script. The power of the AI based model is to visualize the consequence of that change and to use that information to decide to go further down that route, or explore something else.

[292] On Fabula see the recent article by Mirowski et al, "Building a Narrative Storytelling Sidekick with the Writers' Community", arXiv

Alice: Thank you, that's really very interesting. So we can say that Fabula simulates in some way human imagination?

Piotr: I would say it simulates what happens in the world of the story. It doesn't simulate human imagination, because human imagination, I believe, is something very difficult to understand and to simulate, but rather it exploits the capacity of linguistic structures to simulate potential outcomes, in a purely linguistic sense.

Alice: So it does not have imagination itself, but it can stimulate human imagination?

Piotr: I would rather say this way, exactly as you said, that the tool is a visualization tool.

Alice: I like the idea of a visualization that is not visual but it's a linguistic visualization. So it's the visualization of all possible directions that a story could take. I think this is really interesting.

Piotr: Yes.

Alice: When you showed me Fabula, what struck me the most was that as soon as you change a little part of the architecture of the story, which is really a complex architecture, you see all the other parts that start changing. I found this very very fascinating. So my question would be: could it be that the AI could help in keeping our stories open? In the sense that we have a narration and we know how it ends, but we can use the AI to re-open that end, and also the development, and in some way improvise our story? Is there perhaps a connection with the extraordinary work that you are doing within *Improbotics*[293]?

Piotr: Yes. I can give you an example of a workshop that I run recently with young adults at the National Youth Theatre in London. They used Fabula with *Romeo and Juliet*, reimagining alternative versions of the play by changing settings, relationships, or character dynamics—for instance, swapping genders or introducing new conflicts. Fabula generated alternative versions of the script while maintaining internal consistency. The participants then performed these versions through readings and staging. The goal was not to produce a final text, but to enable actors to explore the original play through its alternatives. The real outcome was the performance itself, not the generated script.

[293] On this extraordinary theatre project, carried by Mirowski since a few years, see at least the website: https://improbotics.org/ Note that this project is something that Piotr does outside of his research scientist work at Google DeepMind.

Alice: This is very interesting. I love the example also because exactly two weeks ago I was in Verona and everything there is about *Romeo and Juliet*. And of course when you go there and you see all this, you feel the need to let the story become alive again, because it became a stereotype. The example that you gave is very interesting, because these kids managed in some way to try to let the story become living again for them. Actually this was going to be my third question. Very often people worry (and they are right) that the AI produces lots and lots of clichés and stereotypes and biases. This is due, of course, to its own way of functioning, as the AI learns the recurrent structures of data. So my following question was going to be this one: what strategies can we adopt so that this aspect of AI turns into its opposite? I think you are doing this in an extraordinary way through *Improbotics*, but when we use Fabula, perhaps we can do that too.

Piotr: I'm thinking about the tool Fabula and the current usage of large language models in the context of programming. That is effectively the biggest area of development and excitement around large language models: the fact that they can be combined with agents to actually write code and to almost write it themselves. Large language models have two potentials which can be useful to some extent. One is the generation of alternatives: the fact that they are stochastic means, that every time you generate something, it might be slightly different. It's important to see large language models not so much as oracles who give a definite answer, but rather as a way to sample potential outcomes based on a context. Fabula is presenting with the options to actually sample again and again, retry, rerun, to see what alternative options there are.
I believe that is a mindset that we need to always keep in mind, to approach AI not as a source of truth but rather as a source of untruth, which has to be approached with a lot of critical thinking and handled with care, as something that allows us to quickly generate alternatives and explore those alternatives. The second advantage of large language models is that they have been trained to follow instructions and they became relatively good at following multiple instructions at the same time: multiple constraints. So in the case of Fabula the language model is following hundreds of instructions at the same time. For instance, saying that this character has this appearance and these intentions and that character has this background and those intent and these intentions, and that this happens at that part of a scene in the first beat and then in the second beat, and the location has this appearance – you combine all those elements together, all those instructions, to ask the language model to try to take all those instructions at once to produce an outcome. And you're effectively being in the position of a programmer, who is manipulating the machine by changing the knobs and the various parameters which are exposed to you as you use that machine. So when you use Fabula in the architect mode, you see all those elements of the structure of a story, including questions such as: what is the suspense? What is the existential question? What are the relationships between the characters? And

you can modify them one by one. You have effectively a lot of control over the outcome produced by the language model.

Alice: Thank you Piotr, this is a very interesting aspect. Perhaps many very interesting stories also in literature were born this way: what if we changed one element of a very known story? What would happen? Let me ask you one more question. What motivated you to work on Fabula in the first place?

Piotr: I started to work on a first version of the tool in collaboration with Richard Evans and Kory Mathewson around 2021–2022. The tool was called Dramatron, and it provided a simple scaffolding of stories by decomposing the generation into a logline, a list of characters and descriptions of characters, a list of plot points, location descriptions, and finally dialogue for each individual scene. We had run extensive human-computer interaction studies with playwrights and screenwriters using Dramatron, and we published our findings in 2023, a few months after ChatGPT was released. But we had worked on our tool completely independently and before that happened. Richard continued working on narratology and storytelling and had a first prototype of a drama manager within Fabula, I think in 2024. He shared with me some examples, and I was for the first time surprised by the output of a language model, something I hadn't been for a couple of years. That gave me the impression that the language model was relatively good at simulating what happens in the universe of a story. The prototype, called "Dramabox" was keeping track of the character intentions and it had a narrative arc for the characters. And there was also a narrative arc in the story that would be dynamically changed and adjusted as the story progresses and as my instructions would change. As a result, Dramabox would not only directly respond to what I would say and try to shift the story in that direction, but would actually continue along some trajectory that did make sense. I was positively surprised that it did. I felt there was some level of agency from the machine in how the characters were evolving in the story.

Alice: Wow, this is very interesting and also, from a certain point of view, quite disquieting. One of the questions that my students always raise is: does the AI understand? Does the AI have consciousness or so? Probably you would never say at the moment that it has consciousness or that it understands. But this fact that you're saying, that it goes on through a path and continues inventing the story: does it mean that in some way it has found a kind of orientation?

Piotr: No, the reason why I was surprised is because I didn't see the structure of the narration: because it was hidden from me. I only saw the outcome of the text, but effectively what the tool had done was that it had planned a narrative arc for the characters and a new narrative arc for the story. And it was going along that plan. What I was testing at that moment was a slightly different version of Fabula, Dramabox, which was more like a "choose your own adventure game", where I would constantly respond to

what happens and the tool would generate something in response to my instructions. But I felt surprise because I was seeing only the external aspects of Dramabox's outputs. I didn't see the internal workings. There was nothing magical and nothing autonomous or independent behind it. It was simply the power of a language model producing a plan and following that plan, and then dynamically changing the plan based on my inputs and what happened in the story. That created an illusion of agency.

Alice: Okay, I see what you mean. So you're not saying that it had an agency, but that there was an illusion of the agency because you didn't see the structure of it. While through Fabula, you mean, the structure becomes evident. Does the whole structure become evident, or is there still something that you cannot really see but is happening?

Piotr: No, you effectively see the entire structure of Fabula. However, you don't need to constantly watch it. If you want to explore the story in some directions, you can also switch to what we call "gardener mode" and see what happens if you write: "In the next bit, this is what should happen". You see how the tool takes your instruction and all the story it has in its structure to produce that bit of paragraph in response to your prompt, but also in response to the entire structure that's behind it. This is one of the criteria for a computationally creative system, that it manages to surprise the creator. And I felt that by effectively hiding the structure, you could get that element of surprise. But also, I really wanted to insist on the fact that the reason why AI tools seem so agentic and autonomous today is mostly because we don't necessarily see everything that's going inside, but rather we only see the output. And we ascribe an interpretation of agency to those tools.

Alice: That's very interesting. I could ask you a question about this: on the one hand you are saying that these systems seem agents because we don't see what's going on in the structure, but on the other hand you were saying that they seem creative or they help us to be creative when we don't see what's going on inside. So, is it better if we see all what's going on inside, so that we can take away this impression of the AI being something magic, or is it better if we leave some of this black box structure, because this helps us to interact with it in a creative way?

Piotr: Both. These are exactly the two modes and assumptions that we explored when building Fabula. We had an "architect mode" that lets the writer see the structure and then of course removes all the amount of surprise, because all the instructions are there, visible, and the writer can navigate through those instructions. And alternatively there is another mode called the "gardener mode", in reference to a saying by George RR Martin, the writer of *A Game of Thrones*, who said there were two kinds of writers: the gardeners and architects. The gardeners ride in the flow, whereas the architects want to plan ahead. In that second mode, the gardener mode, because we're hiding information, there is more potential to be surprised by the AI, and maybe also ascribing creativity to the machine. I believe that it can be useful to some extent for our own creativity to suspend our disbelief

and to go beyond the analysis of what happened, into the imagination of what could be done with that new bit of information. This is essentially what we have been doing on stage, when doing improvisational comedy with robots and language models. We are surprised by the outputs coming from a machine and we decide not to criticize them but rather to take them at face value to see where this leads us in our creative process.

Alice: Oh yes, the *Improbotics* shows are incredible from this perspective. For readers who didn't see any of them yet: there is an AI on the stage, whose name is Alex, and both actors and public improvise on the outputs of Alex. This is very interesting because in some way they take the responsibility of these outputs. They don't say: okay, this was said by an AI, so it's meaningless, but they take it and transform it in their own way. Something similar happens, you were saying, also in Fabula through the gardener's modality.

Piotr: Yes, that was the intention when the tool was being built. I have to say I haven't used that mode that much, simply because I was testing and building the architect mode (and seeing the entire structure of a story in front of me allows me to easily test and debug if it's working) and because the writers with whom I worked were always trying to rework their existing material. So they were happy to actually see the analysis by the machine. But definitely the garden mode is a sort of immersive mode that could be explored by writers.

Alice: Maybe we could try, when we start to use Fabula in our course, to use both modes, to see what comes out using them[294].

Piotr: Yes. And hopefully I'll be interested to see if there is a feeling from the users in the gardener mode that the story is not constantly fighting them: that this system is not fighting them trying to impose its own vision.

Alice: What do you mean?

Piotr: One of the limitations of the gardener mode, which is something we observed when writers worked with it, was that they sometimes felt it was not very flexible. It continued along its own direction.

Alice: Now I understand. We experienced that too, when we tried to do some visual experiments with the students. We were asking for a particular image and the AI continued to propose a certain image. The students became very annoyed, because they couldn't convince the AI to leave what it was doing and start something different.

[294] Fabula will be used in the course "Aesthetics of New Media," taught by Alice Barale and Maddalena Mazzocut-Mis at the University of Milan. The results will be presented at the conference "Making Sense of Slop," which will take place on May 12, 2026, at the Department of Cultural Heritage.

Piotr: Yes, I think controllability of AI tools is very important and those tools are controllable in different ways than human creative partners. Quite often we use the same sort of logic of trying to convince them as we would do with a human partner. We are frustrated because of course the tool will not react and respond in the same way.

Alice: This is a very fascinating aspect. Prompting is so strange, because it really seems to follow its own rules, like human exchange sometimes.

Piotr: Yes. And what's frustrating is the lack of control. The whole fact that there was a prompt engineer who was inserting specific keywords in the prompt highlights the fact that these are relatively strange devices to get any task done. And I think most of the research recently has been exactly on trying to make the prompting more consistent and explainable.

Alice: Thank you Piotr. It's very interesting also from a philosophical point of view, because you know, philosophers are reflecting on creativity since hundreds of years and what you're saying basically is that you have to keep together the two aspects: explainability and surprise. To keep these two aspects together maybe is the big challenge, even when you don't use AI.

Piotr: I fully agree. Yes, this is this is the key problem that is constantly being addressed by anyone trying to use the tools for creative purposes.

Alice: Thanks Piotr, I think we have I have understood a little bit more and we are ready to start experimenting with Fabula. We would be very happy if have some outcomes to share in the near future, people experimenting with Fabula who want to talk more about it.

Piotr: Yes, there is a film director, Martin Percy, who is using Fabula to create an interactive story that has multiple ramifications – he has been using Fabula to explore different paths in the story – and we have a paper going out soon on Arxiv has a deep analysis of the interviews. with Fabula users

Alice: We are looking forward to it. Thanks for the very interesting conversation.

295