Universiteit van Amsterdam

Max Impact

The Making of Affiliaterobot.nl

June 2, 2023

Student:
David Werkhoven - 13909495
Mette Lijmer - 13983814
Chris de Hart - 13937316
Frenklin Meijvogel - 14007037
Jasper van der Valk - 13854577
Rick van der Veen - 13984330

TA:
Dionne Gantzert

Course:
Tweedejaarsproject

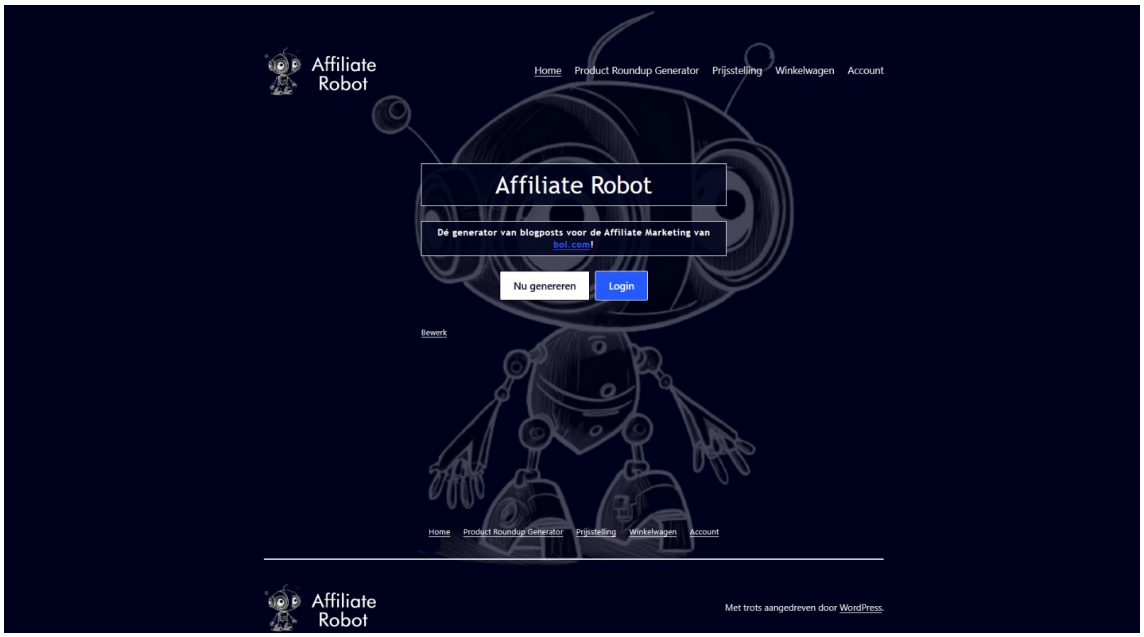

Contents

1 Introduction

AI has seen many new and exciting developments in the last few months, the biggest of these being ChatGPT. With this new way to generate text or conversation, many other AI tools came along like Koala.sh. Koala.sh is a website that builds upon ChatGPT. One section of the website, called KoalaWriter, is where users can generate articles. The website does this by using the OpenAI API. One of the article types that users can generate is roundup articles of products from Amazon. The affiliate marketing link of the users will then automatically be added to the product links in the article. KoalaWriter is thus a convenient way for affiliate marketers from Amazon to create articles. Bol.com is more famous in the Netherlands than Amazon. Bol.com also has an affiliate marketing program, so the company Max Impact wants to use this untapped market to introduce this platform here as well. The company Max Impact is a one-man-start-up company that would like to see a Dutch version of KoalaWriter with Bol.com. The company would like to build an article writer, which is able to generate roundup articles with the help of ChatGPT, just like KoalaWriter. The articles should include a variable amount of products of a certain type, for example sunscreen or dog food, and showcase each product with a corresponding product description and a list of pros and cons. The articles should include a link to the corresponding products that can be found on the Dutch commercial platform Bol.com. The links should contain the user's Bol.com affiliate code, which gets the user a small payment for every purchase done using that link. To display the products in the article, web scraping will need to be used to fetch the best products off of Bol.com. To use ChatGPT for the articles, the OpenAI API is needed. The price of using this API is covered by Max Impact. The end goal that was discussed is a web application, where users can pay to have the website generate articles, which in turn, can generate revenue for the user using the Bol.com affiliate program. The website should be designed for users who wish to spread articles about products on for example their own blog website. However, due to the affiliate links, the users can spread their articles however they wish.

To generate an article, the user should be able to specify a certain set of keywords. For example, the strings 'dog food' or 'sunscreen' should generate valid articles. The user should also be able to specify a set of parameters, which influence how the articles are written. These parameters include the ChatGPT model used, the amount of products that need to be ranked and the tone of voice and point of view that the article should be written in. Design wise, the user interface of the website should be clean and easy to use. A login system and a way for users to add their Bol.com affiliate code should both be present. Finally, a payment system has to be implemented in order for users to buy tokens to generate articles, which in turns generates profit for Max Impact. To achieve our goal of creating this product, we have split the work into 3 major parts. These being front end, back end and prompting ChatGPT by using the API. We will accomplish our goal by completing multiple tasks. We need to show the company different designs for the front end, and work on the one that is chosen. The website should be able to handle users spending money to buy tokens which in turn can be used to generate their articles. We will work with several plugins. Some will take care of the payment and account system and one will have to be built by us to handle the prompting of the API. By doing all these things we will have met the wishes of the company.

In this report, we will provide a summary of our progress which leads to how we built the final product. In section 2, we will explain the demands of Max Impact. In section 3, we will discuss our implementation and the problems we faced, together with our proposed solutions. We will do this in three parts, namely the API usage, the front end and the back end. In section 4, we show the finished project with some screenshots. Finally, in section 5, we will summarize our product and discuss some limitations and ideas for future improvement.

2 Demands

Before we got started coding our implementation there were already some important decisions to be made, as Max Impact had multiple, different demands for the final product. First of all, the website needed to be designed in the WordPress[1] environment, seeing as the company deemed themselves proficient in editing websites via this software. This way the company can work on further improving the website after the project has finished, without hiring new people. This demand is not one that comes easily, as to work with the ChatGPT API in WordPress a separate plugin must be created. Designing a solution with Node.js[2] was also an idea, seeing as our group already had some understanding of this runtime environment. But after a discussion with Max Impact, which concluded with the fact that the company could not maintain this project efficiently in Node.js after the project, we decided to work with WordPress and to subsequently create our own plugin.

A second important part of the product is a working payment system, through which users can buy tokens. We decided on the plugin WooCommerce. WooCommerce is a popular WordPress plugin, famous for turning a WordPress website into an online store. It has many extension plugins, such as ones that allow different types of payment. The most important extension allows the implementation of subscriptions as products. Since the main goal of Max Impact was paid article generation based on subscriptions, this extension plugin helps to accomplish that goal.

The third part that is crucial to the project is web scraping, seeing as the generated articles need to use information found on Bol.com to write a fitting paragraph. The web scraper makes use of the search engine of Bol.com to list the most relevant products on Bol.com based on the keyword that is given by the user. The top number of products from the search result will be used. For every one of these selected products the price and the URL will be taken. The price will be saved and the URL is used to get the page of the product. From this product page more data will then be retrieved. This data includes the name of the product, the main image, the product description, the specifications and the average review score, a rating between 1 and 5. The last thing retrieved are the reviews of the product, including the pros and cons. This all results into a batch of useful information which can then be passed to a script that includes the OpenAI API to generate the necessary text. Finally, the domain name we had to settle on was decided. After a discussion with Max Impact, the company decided on the domain: Affiliaterobot.nl. The company then hosted the website for us to edit.

[1] WordPress is a popular and user-friendly website creation tool that allows users to build and manage websites without needing to write code, making it accessible for beginners.

[2] Node.js is a powerful and versatile runtime environment that allows users to run JavaScript code on the server-side, enabling them to build scalable and efficient web applications.

3 Method

3.1 The implementation

Building a website is a complicated task. While working on this project, we encountered numerous obstacles that we had to overcome. The notable problems and their corresponding solutions will be presented in this section. We also separated the task into 3 major parts: front end, back end, and API usage. These are all individually explained and expanded upon below.

3.1.1 The API

The API key for the gpt-3.5 and gpt-4 models is provided by Max Impact. The main language of choice to use the API and engineer the prompt was Python. The task can be simplified by splitting the article into multiple parts; the introduction, the product paragraphs and the conclusion. The product paragraphs each consist of a pros and cons prompt and a description prompt. This allows us to generate multiple different paragraphs instead of one big article, to insure a more targeted approach for each part.

First, the introduction is generated. When starting with a simple prompt, little quirks are often encountered like explicitly repeating the input query back in the introduction. To avoid this, little alterations were made to assemble a prompt piece by piece which generates a proper article introduction. Then the product parts are generated. According to the wishes of Max Impact each paragraph should cover a product with a list of 3 pros and 3 cons of the product followed by a description. Due to the web scraping, a lot of product context is available from Bol.com, so we had to decide which information was relevant enough to include in these prompts. Information like the reviews and the name of the products are crucial to writing this part of the article. Data such as the price can be useful to include as extra information while we chose to omit context like weight or product specific details, such as the amount of pages for a book. The relevant data is then extracted from the results of the web scraping and put into the prompt as context. There were some small problems that needed to be solved, like the generated text always mentioning the exact name of a product, which is often way too long, or mentioning that it was in fact going to list pros and cons. We also added some strict guidelines that the generated text should follow, like having the pros and cons formatted in a way that satisfies the company. After this was completed, we had constructed the prompts that would generate the paragraphs according to our wishes. Repeating these prompts the same amount of times as the amount of products with the appropriate context for each product results in the middle paragraph of the article. Finally, the conclusion is generated. To write this conclusion, we pass all previously generated text into a prompt with the added context of having to write a suitable conclusion to the article. This all concludes the method of actually generating the article. Seeing as this was done in Python, a translation had to be made to JavaScript so that it could be used on the website.

To evaluate the prompts used, a couple different methods were used to improve and check the usefulness of the prompts. When engineering a prompt, some obvious discrepancies can stand out. A lot of products on Bol.com have a really long product title, for example with specifications in the title. Asking for a paragraph about a product with a very long product name can result in ChatGPT using the entire product name in the paragraph, which takes away from the readability of the paragraph. Another small issue was ChatGPT explicitly mentioning in the article that there would be a list of pros and cons in the next paragraph. These issues could be easily fixed by making the prompt more specific and including more detailed instructions. We also had direct contact with the company to insure the quality of the articles was up to standard. Comparing our generated articles with the articles that are generated from KoalaWriter was another helpful way to evaluate our prompts.

3.1.2 The front end

Two possible front end designs were made in WordPress, both to get an idea of how to use the WordPress environment and to give the company a choice of layout. After designing both the front ends, Max Impact decided on one of our designs, so after this decision was made, we could focus adding the necessary features. There are 5 main pages on the website. The first page is the home page, with a button leading to the generate page and a button leading to the login page. The second page is the generate page where a user can generate product roundup articles, shown in Figure 1. Here, a user can fill in multiple fields to change the output article. These fields include: buttons to change the model between ChatGPT 3.5 and ChatGPT 4, an input box to insert the desired product to use for the roundup, the amount of products to rank for the roundup and the tone of voice and the point of view that the article should be written in. After filling in these fields, the user can click the generate button and the article will be generated. The user can see the generated text being typed in real-time, and can cancel the generation at any moment. The third page is the pricing page, where the user can put tokens in the cart. The fourth page is the cart page, which is where the tokens can be bought to use for the generation. The fifth and final main page is the account page, where the user can register, log in and log out. After logging in, the user can change their account details or their affiliate link. The theme and colours of the website are inspired by KoalaWriter. A logo was also added to add to the overall layout of the website, while also creating an icon that makes the website more memorable and recognisable to users which makes them more likely to return.

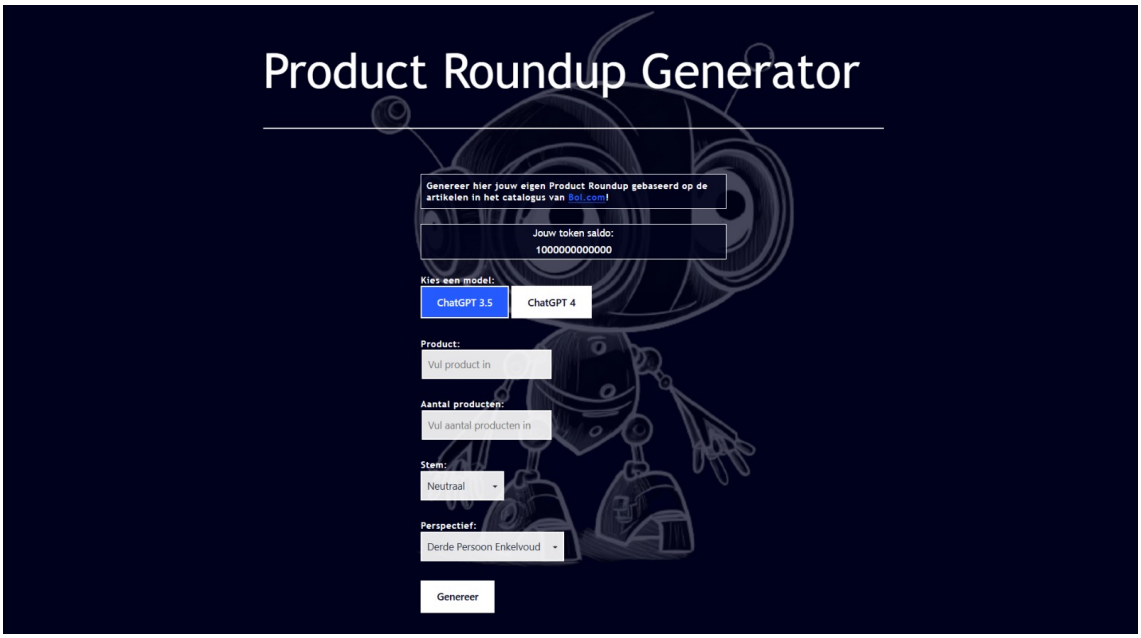

Figure 1: The generate page

Most of the functional aspects of the front end originate from the WooCommerce plugin. This plugin automatically adds different pages to a website, such as a shopping page, a checkout page and an account page. The plugin also allows users to make and customise different types of products. A user can add as many products as they like, but we chose to have six different subscriptions tiers where each subscription gives the user a certain amount of tokens per month. Besides regulating products and payments, the WooCommerce plugin has more functionalities. One of them is a premade account system that includes features like changing user's passwords and showing completed orders. However, the WooCommerce plugin lacked a part that is essential to the website, this being an input box for a person's affiliate link. This link is used as part of the hyperlinks inside of the generated pictures and titles of the products in the articles. We made a separate plugin that uses the WooCommerce API resulting in the account system being changed slightly by adding a savable input for the affiliate link.

To prevent users using the generator while not being logged in, a plugin called 'Ultimate Member' is used. This allows the admin the redirect a user trying to go to a specific page towards a different page when not logged in. When trying to access the page containing the generator, they will be redirected

towards the login screen that has been made using the same plugin.

3.1.3 The back end

The back end consists of several important parts. Firstly, we had to create a plugin which would be supported by WordPress. The purpose of this plugin is to prompt the OpenAI API with information which has to be gathered from Bol.com to return a human-readable product roundup article. The back end is the backbone of the plugin. It connects the OpenAI API requests and responses with the front end where it is processed to be human-readable text and is displayed on the website. One of the main hurdles we had to overcome was the fact that using JavaScript, while easier, created a security vulnerability by not hiding the API key for potentially malicious users. Due to this risk, the choice was made to code the plugin in the coding language PHP. Although this was more complex, the WordPress environment requires this approach to hide the API key on the server side. This makes sure that no users can reach it.

Other developments that were made within the plugin in PHP were a shortcode and a stream with the OpenAI API. Out of these ideas the stream was the most challenging since it required high knowledge of PHP. The OpenAI API makes it possible to send a response of a prompt via a stream. This means that parts of the response will be send immediately the moment they are generated. This stream made it possible to display an article live while it generates the content. This way there is no delay while waiting on the article to appear and spotting a disruption becomes easier. Implementing this effect was also a demand by Max Impact. The shortcode was needed to dynamically alter and update the website's content. The prompting requires a lot of input boxes such as the amount of products or the product type. These input forms were generated with the use of a shortcode. This API implementation handles the introduction, middle paragraphs and conclusion separately and had to be implemented.

Apart from the generation part of the back end, a token system was implemented that would result in a change in a users token balance from buying a product and generating an article. Tokens are the currency of Affiliaterobot.nl that can be bought in the shop as either subscriptions or bundles and are needed for a user to be able to generate their own articles. Thanks to the WooCommerce plugin, a shop and account system was already made. The subscriptions, bundles and shop page can be seen in Figure 6 in the Appendix. The token system adds the token balance to the database that rises when the tokens are bought and depletes when an article generates. Whenever the token balance is too low, it will not allow the user to generate an article.

If an article is generated, the product titles and pictures contain hyperlinks towards the users affiliate version of the respective products. This is implemented in the back end as part of the account system by saving another entry in the database using a hook (or API) from WooCommerce. That means that it is an extension for the WooCommerce plugin instead of a plugin that works on its own.

4 Result

In this section we will show the final results. In accordance with the company, we have chosen to keep the code in a private repository. The website Affiliaterobot.nl is however a public website which can be viewed by anyone. When a user visits the website, they will first arrive at the home page, shown in Figure 2. This page shows the user what the website Affiliaterobot.nl is about. They can click buttons to generate, login or go to any other pages shown in the menu at the top-right of the page.

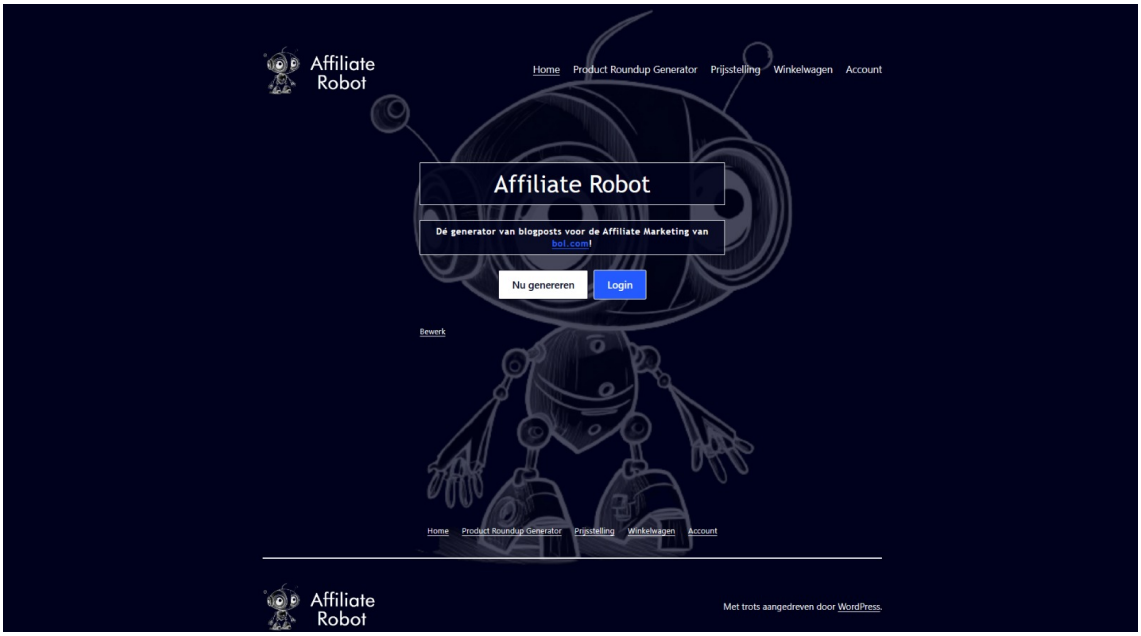

Figure 2: The home page

If the user clicks the generate button they will be sent to the generate page, shown previously in Figure 1. This is the page where product roundup articles can be generated. Before generating the user has to fill in the different input forms which will change the parameters of the generation. The user can choose between using ChatGPT 3.5 and ChatGPT 4 for the generation, with ChatGPT 4 being the better but more expensive option. Then the user has to fill in the type of product and the amount of products they want to rank. Afterwards they have to choose the tone of voice and the point of view that the article will be written in, using a drop-down menu. The final step for the user is to click the generate button, which will begin generating the article. The article that is generated starts with an introduction. Afterwards it has a paragraph for each product, with a short of description and a list of pros and cons. The article ends with a conclusion. Also included are the price and average score of the products, which are situated underneath the image. The title and image of each product are clickable and, using the user's affiliate link, lead to the corresponding product page on Bol.com. Figure 3 in the Appendix shows an example of one of the product paragraph of a generated article about dog food. Lastly we found out that, while the streaming of the generation did work on our localhost servers, it didn't work on the final server. This was due to the fact the final server was a shared WordPress server which didn't allow streaming. We found this out relatively late, but still decided it was important. As such, we created a no-stream version of the generation plugin, which works perfectly on the final server.

In order to generate the previously mentioned articles, the user needs to pay tokens which can be bought on the website. In order to keep track of this, the user needs to be logged in. This requires the account page, together with the login and register pages. The login page can be seen in Figure 4 in the Appendix. Once logged in, the user can access the account page, where they can change their account data, including the affiliate link. This page is shown in Figure 5 of the Appendix. The tokens can be placed in the cart on the pricing page On the website. There are different options to choose from, which are seen in Figure 6 in the Appendix. Once placed in the cart, the user is redirected to the cart page, where they can purchase the tokens using different payment methods. The cart page can be seen in Figure 7 in the Appendix.

5 Conclusion

To conclude, we created a WordPress website for Max Impact called Affiliaterobot.nl. The website uses the OpenAI API and a custom plugin to generate custom-made product roundup articles, which rank a certain type of product found on the website of Bol.com. The generation takes a fee in the form of tokens, which can be purchased on the website. Each product can be clicked and leads to the corresponding product page on Bol.com. Users can also link their own Bol.com affiliate link so that when the articles are shared across the internet, they earn a small commission through the Bol.com Affiliate Program.

5.1 Discussion & Future Improvements

Aside from the successful implementation of Affiliaterobot, the website also has some limitations which will be mentioned in this section. The future improvements of these limitations will also be mentioned. Firstly, the mobile version of the website is not as appealing as the desktop version, since we didn't focus on the mobile version. Further improvements should focus on creating a separate layout for this version. Furthermore, we only tested the website on the Microsoft Edge and Google Chrome web browsers, meaning that using other web browsers might cause errors. Thirdly, the website does not provide any contact information, meaning that users cannot file a bug report. This should also be added as a further improvement. The optimal pricing of the tokens should also be researched further, in order to maximize profit for Max Impact. Currently the pricing is equal to the pricing of KoalaWriter. To add to this, the term tokens can be vague, as it is not synonymous with words. One token does not equal one word. To avoid potential confusion, this could be clarified.

The functionality of the website could be extending by adding more article types to generate. Examples of these article types are: blog posts, in-depth reviews or text summaries. This would be easy to add, as the framework for generating text based on input and processing it is already present in the back end. Early on we discussed an idea with the company to embrace AI technology even more and include artificially generated images. We dropped this idea as we and the company realized this addition fell outside the scope of the original product idea. This could however be added as a future improvement. Furthermore, the prompts could be optimised further in order to improve the output quality of the generated articles. Penultimately, it was a shame that the streaming didn't work on the final server. As a future improvement, Max Impact could switch to a private server that does support streaming. To make this as easy as possible, we have created separate plugins in the shared Github repository so Max Impact can easily switch between versions. Finally, the most important future improvement is the addition of a history of generated articles. This feature is already present in systems like ChatGPT and makes sure that generated text is not lost when leaving the page. Even though several warnings are given when a user attempts to leave a page when an article is generated, being able to save the generated text would increase the quality of life of the website.

5.2 Final remarks

In short, we worked together with Max Impact to create the website Affiliaterobot.nl. After a successful collaboration, we have fulfilled the company's wishes and as such have put the company in a position where they can share the website and expand upon it.

A Appendix: Additional Website Screenshots

Figure 3: An example of a generated article

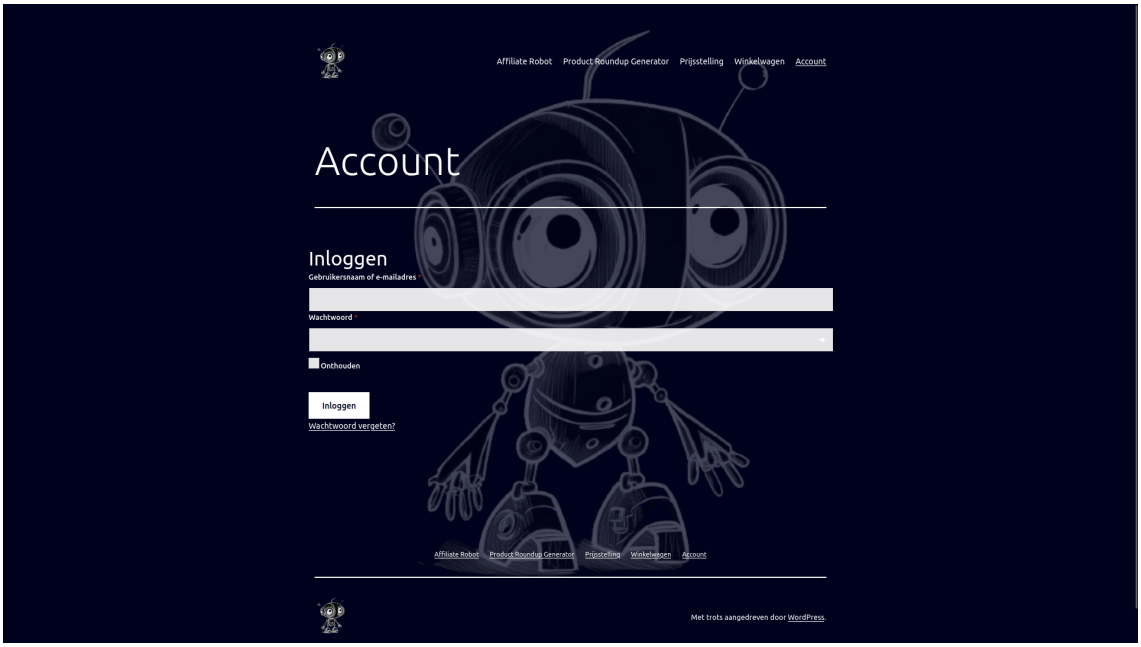

Figure 4: The login page

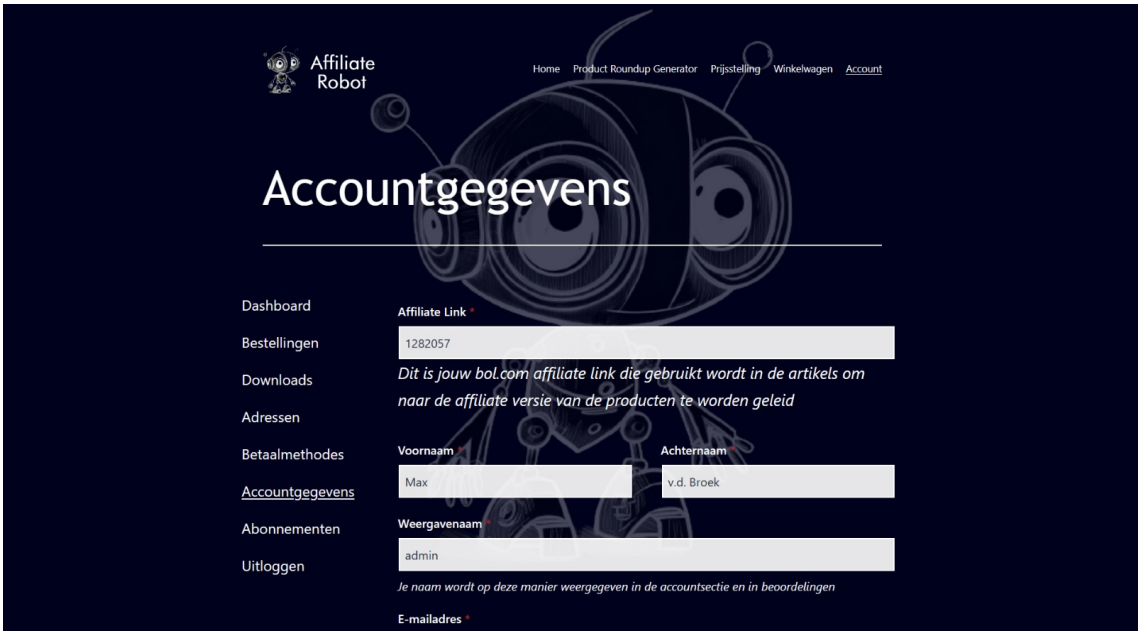

Figure 5: The account page

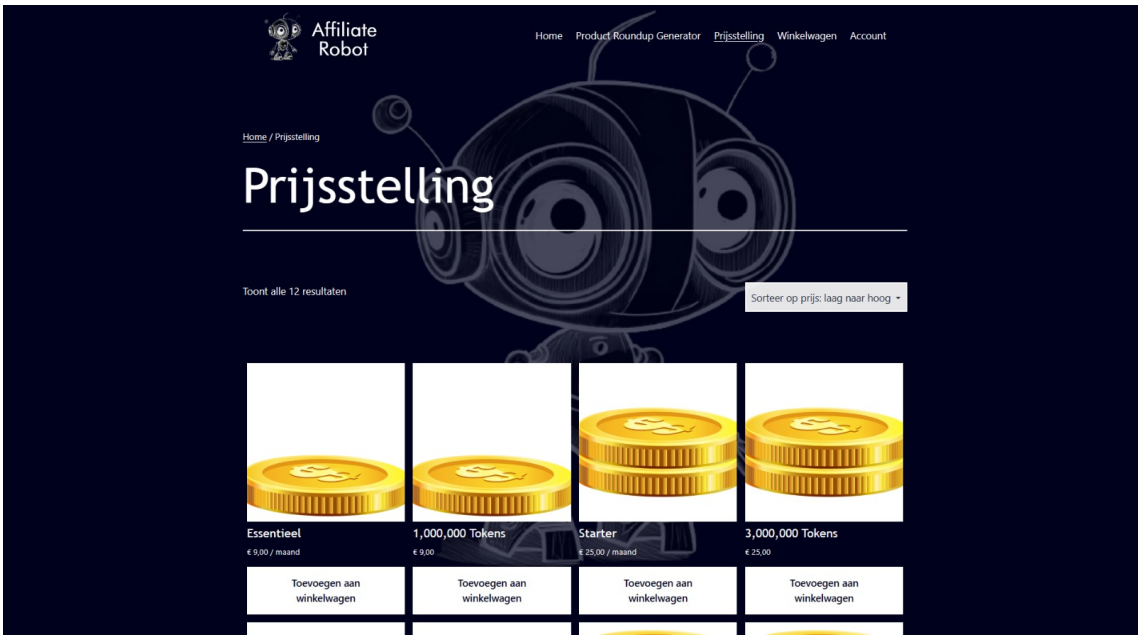

Figure 6: The pricing page

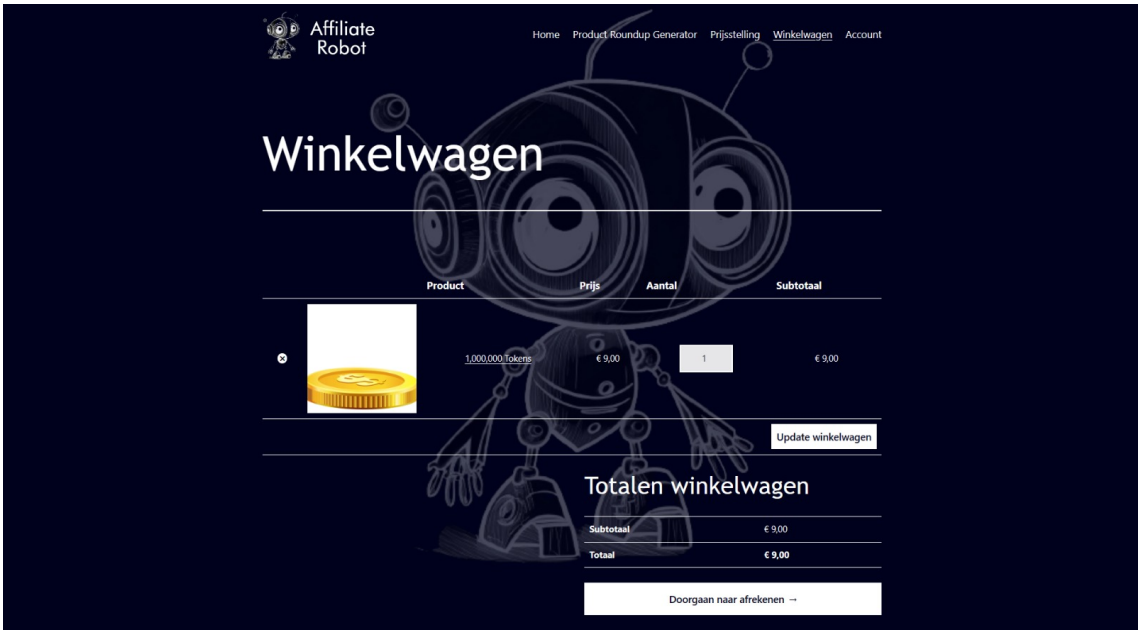

Figure 7: The cart page