

AI FOR CONTENT AUTHORIZING



GREG DUNLAP / CONTENT STRATEGY SEATTLE / FEBRUARY 2026

GREG DUNLAP

Building and supporting CMS implementations
for 20 years

Architect of Drupal's configuration
management system

Internationally ranked pinball player

<https://www.linkedin.com/in/gdunlap/>
<https://bsky.app/profile/heyrocker.com>

Designing

CONTENT AUTHORING

E X P E R I E N C E S

By Greg Dunlap

I also wrote this book *Designing Content Authoring Experiences* which I self-published last year, which many people in this room supported and I am super grateful to everyone for that. One of the ideas that I had when I wrote this book was that a lot of the discourse I was seeing in the CMS world wasn't relating back to the clients I was working with. They weren't worried about personalization or orchestration or AI because they were failing at the most basic elements of launching a successful site, they were just very immature and they were not set up for success from an organizational or governance standpoint. So I made a book that was very approachable and kind of goes back to basic principles. If that sounds interesting to you, I'd encourage you to check it out.

In my experience, the number one complaint of organizations looking to replace their CMS is, simply, “our editors hate it.”

Deane Barker - Real World Content Modeling

So I like to start all my presentations with this quote and I'm sure many of you have seen me use it before. But I would argue that most authors don't hate their CMS, what they hate is systems that are not designed with their needs and use cases in mind. They hate being disenfranchised from the very system they are expected to use day in and day out. And this is really important because when you spent millions of dollars building an enterprise CMS, you want people to actually engage with it. I was at a conference recently where someone said their editors do 6-9 THOUSAND edits a *week*. That is a lot of people doing a lot of work in your editorial tools. If they are filled with dread every time they log in, then guess what? They won't unless it is absolutely necessary, and once they are in they will try to get out as quickly as humanly possible, and this is really why I wrote my book. Because authoring is important. It prevents people from being miserable *and* its good business.

So this talk is about AI and I've historically been kind of a skeptic about AI, especially in the realm of content authoring. Everything I was seeing was focused on generating content faster to feed the maw of consumption more efficiently, and that is not a use case I've been very excited about. However in the last six months or so I've started to come around and see the ways that AI can be used as a tool to *enhance* content production rather than replace it. And thats really what I want to focus on in this talk. That said, there is something I want to get into first and its going to seem like a tangent but its really not (or maybe it is and I just wanted an excuse to have a little rant, and if so sorry.)

Hot Topics SEO

The Riveting History of SEO: From the '90s to Today

Last Updated on January 12, 2026 | Published: December 20, 2023



Ready To Transform Your Content Marketing?

Launch your first Relliverse AI content strategy agent for \$499. Our starter package includes a consulting session with a Rellify expert. Begin producing high-quality content optimized for AI search in minutes!

[Get A Demo](#)

Table of Contents

1. When did SEO begin?
2. Who invented the term SEO?

In the last 25 years some really interesting things have ben happening on the web. In 1997 the term SEO was coined.

Web Content Accessibility Guidelines 1.0

W3C Recommendation 5-May-1999
superseded 18 May 2021

This version:

<http://www.w3.org/TR/1999/WAI-WEBCONTENT-19990505>
([plain text](#), [PostScript](#), [PDF](#), [gzip tar file of HTML](#), [zip archive of HTML](#))

Latest version:

<http://www.w3.org/TR/WAI-WEBCONTENT>

Previous version:

<http://www.w3.org/TR/1999/WAI-WEBCONTENT-19990324>

Editors:

Wendy Chisholm, [Trace R & D Center](#), University of Wisconsin -- Madison
Gregg Vanderheiden, [Trace R & D Center](#), University of Wisconsin -- Madison
Ian Jacobs, [W3C](#)

Copyright © 1999 W3C (MIT, INRIA, Keio), All Rights Reserved. W3C [liability](#), [trademark](#), [document use](#) and [software licensing](#) rules apply.

Abstract

These guidelines explain how to make *Web content* accessible to people with disabilities. The guidelines are intended for all *Web content developers* (page authors and site designers) and for developers of *authoring tools*. The primary goal of these guidelines is to promote accessibility. However, following them will also make Web content more available to *all* users, whatever *user agent* they are using (e.g., desktop browser, voice browser, mobile phone, automobile-based personal computer, etc.) or constraints they may be operating under (e.g., noisy surroundings, under- or over-illuminated room, a hands-free environment, etc.). Following these guidelines will also help people find information on the Web more quickly. These guidelines do not discourage content developers from using images, audio, or video, but rather explain how to make multimedia content more accessible to a wide audience.

This is a reference document for accessibility principles and design ideas. Some of the strategies discussed in this document address certain Web internationalization and mobile access concerns. However, this document focuses on accessibility and does not fully address the related concerns of other W3C Activities. Please consult the [W3C Mobile Access Activity home page](#) and the [W3C Internationalization Activity home page](#) for more information.

This document is meant to be stable and therefore does not provide specific information about browser support for different technologies as that information changes rapidly. Instead, the [Web Accessibility Initiative \(WAI\)](#) Web site provides such information (refer to [WAI-UA-SUPPORT](#)).

This document includes an appendix that organizes all of the *checkpoints* by topic and priority. The checkpoints in the appendix link to their definitions in the current document. The topics identified in the appendix include images, multimedia, tables, frames, forms, and scripts. The appendix is available as either a [tabular summary of checkpoints](#) or as a [simple list of checkpoints](#).

A separate document, entitled "Techniques for Web Content Accessibility Guidelines 1.0" ([TECHNIQUES](#)), explains how to implement the checkpoints defined in the current document. The Techniques Document discusses each checkpoint in more detail and provides examples using the Hypertext Markup Language (HTML), Cascading Style Sheets (CSS), Synchronized Multimedia Integration

In 1999 WCAG 1.0 was released, defining the first ever accessibility guidelines for web content.

Responsive Web Design

by [Ethan Marcotte](#) · May 25, 2010

Published in [CSS](#), [Interaction Design](#), [Layout & Grids](#), [Mobile/Multidevice](#), [Responsive Design](#)

“The control which designers know in the print medium, and often desire in the web medium, is simply a function of the limitation of the printed page. We should embrace the fact that the web doesn’t have the same constraints, and design for this flexibility. But first, we must ‘accept the ebb and flow of things.’”

—John Allsopp, *“A Dao of Web Design”*

The English architect Christopher Wren once quipped that his chosen field “aims for Eternity,” and there’s something appealing about that formula: Unlike the web, which often feels like aiming for next week, architecture is a discipline very much defined by its permanence.



Build advanced skills for growing career opportunities. Choose from tracks in content strategy, UX/UI, communication with data, and learning design.

A building’s foundation defines its footprint, which defines its frame, which shapes the facade. Each phase of the architectural process is more immutable, more unchanging than the last. Creative decisions quite literally shape a physical space, defining the way in which people move through its confines for decades or even centuries.

Working on the web, however, is a wholly different matter. Our work is defined by its transience, often refined or replaced within a year or two. Inconsistent window widths, screen resolutions, user preferences, and our users’ installed fonts are but a few of the intangibles we negotiate when we publish our work, and over the years, we’ve become incredibly adept at doing so.

366 Comments



Become a patron

2010 Ethan Marcotte wrote an article for A List Apart introducing the term responsive design to describe the new techniques that allowed content to respond to different screen formats.

Introducing the Knowledge Graph: things, not strings

May 16, 2012 · 6 min read

 **Amit Singhal**
SVP, Engineering

 Share

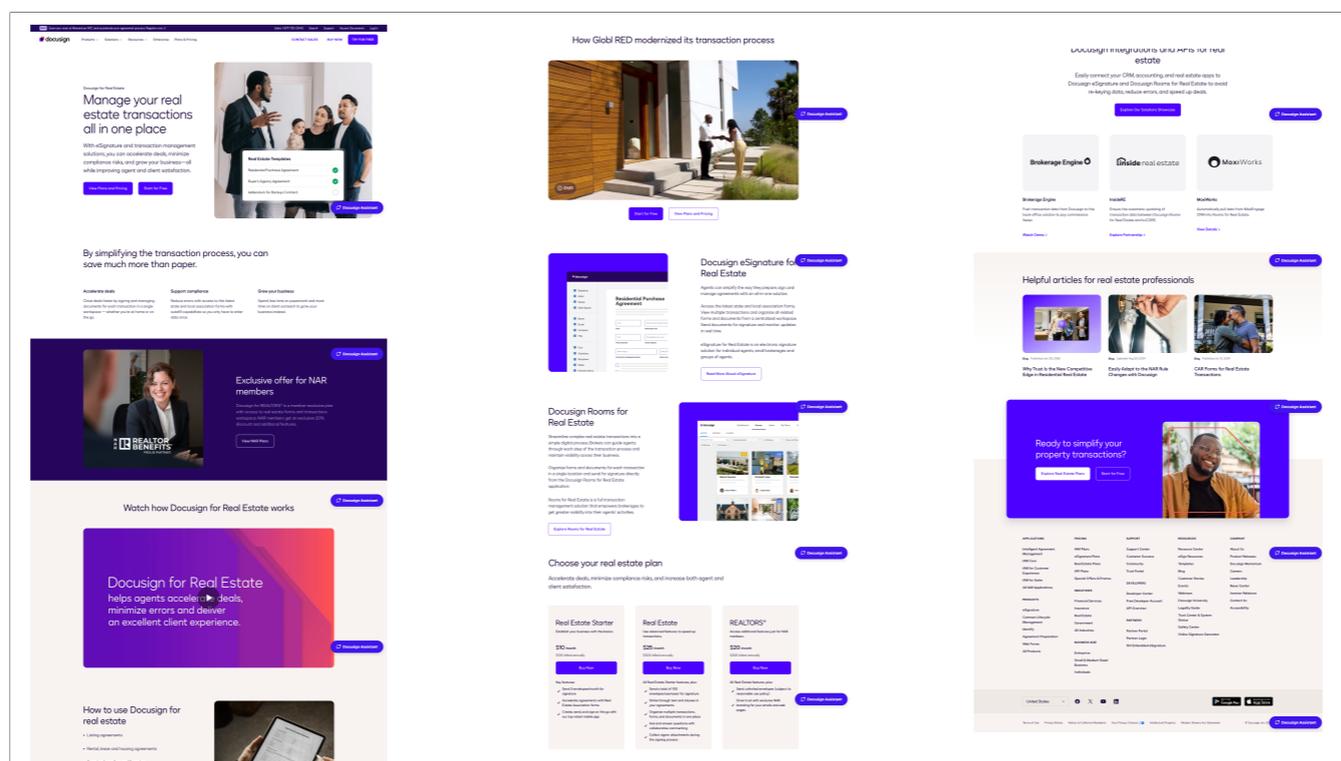
 Listen to article 6 minutes 

Search is a lot about discovery—the basic human need to learn and broaden your horizons. But searching still requires a lot of hard work by you, the user. So today I'm really excited to launch the Knowledge Graph, which will help you discover new information quickly and easily.

In 2012 Google introduced its knowledge graph, a way that content could be described so that Google can know if, for instance, you are describing an event or a location or a tv show or whatever, and then do those special sidebar displays that you see with the map and address and everything.

So what do all these things have in common? They introduced a new audience for the web - and that audience is other computers. SEO is fundamentally about your website telling a search bot what content is on your site and what it is about. Accessibility requires assistive technology to interpret your content based on the metadata provided to it. Responsive design is about describing your content in such a way that your CMS and browser can rearrange it meaningfully depending on the device that is viewing it. The knowledge graph requires that your site describe itself in a specific format that Google can understand and interpret for display. When Amit Singhal wrote the blog post introducing the knowledge graph he titled it “Things Not Strings”, underling this idea that machines need context that isn’t explicit in the text. This is a tension that has been simmering under the web for almost two couple decades now, and which many organizations have dealt with it poorly.

AI is just another extension of this problem. A machine needs to read some content. However instead of removing the need for context, AI actually amplifies it. The knowledge graph requires that you define what something is in order for it to be displayed properly. AI is actually making decisions for you and acting on your behalf. It is generating content and designing digital properties from scratch. I hear a lot of people saying “Oh we’ll just throw all this existing content at AI and it will solve all our problems” when the reality is that without a TON of context about who, what, when, and why then the how has the potential to be much worse than if you had never used AI at all.



So something a lot of organizations have to do is create marketing landing pages. These might be focused specifically on a product, or maybe a seasonal campaign, or a vertical you are targeting, whatever. The basic composition of these pages is a 1-2 paragraph message, followed by some very tightly focused links to content on the site which will support your message. A high profile case study, a white paper, some statistics and comparisons to the competition, and a call to action. These pages typically take a great deal of work and planning to set up, and because of that most CMS companies are offering support to have AI put these pages together for you. This sounds great and I'm sure the pages that it creates will be very pretty however without context the pages will not meet an organization's needs.

So historically, when we talk about things like responsive design or accessibility, the context we want to provide takes the form of structure - creating content types that serve specific purposes, and providing metadata on those content types which provide the context of what the pieces are. Our marketing page example needs this too. If AI is going to go and collect all the content needed to create this page, its going to need to know if its a case study, a product profile, or whatever else. If your content is all dumped into the body field in Wordpress posts, then the AI will have no idea whether something is an event or a press release or a case study or what. It can try to figure it out based on the text, and it will try that, but it will often fail or make mistakes. But this isn't just a problem in our generic Wordpress example. I've seen PLENTY of "structured composable" systems whose components are things like "Promo card with photo" or "2 column layout". These are just as useless because they only describe how something looks, not *what it is*. That promo card could be literally anything, so how is the AI supposed to make smart decisions about the kinds of content it is pulling in to communicate with a customer if it doesn't actually know? So structure to communicate semantic meaning is still just as, if not more, important than ever.

But there's another angle at play here too. A landing page is trying to achieve some goal, some purpose. The AI needs this information too. So the context for AI is not just structural, it is also strategic - What are we selling, what are our goals for this page, who are our customers, what are their needs, how do we plan to move people through the funnel, what does success look like? In the past we've often used things like taxonomy to tie different types of content together on a theme and then pull marketing pages together using that, but given enough context AI has the potential to do this much efficiently and with better results. Now I've seen humans create these

pages without pieces of this context, but they usually have some general framework of an organization and how it works and its broad goals. Even if your org is bad about making these things explicit, over time people start to figure it out. Computers don't have the time and experience to figure it out, you have to tell them. Creating these marketing pages without that context is worse than not having them at all, it has so much potential to be actively harmful to an organization. We have at this point all seen how hallucinations can cause AI to act crazy, is that what you want for your marketing pages?

CONTEXT

In conclusion, the things we're about to talk about are all exciting and interesting, but they don't exist in a vacuum. Systems still need to be designed well. Context is still crucial. Thought and planning needs to be brought into every implementation. Technology projects fail all the time without this, and AI is no different. I know this talk is about authoring, but all this stuff is intextricably related. You can't separate the two.

CONTEXT MANAGEMENT

AI Contexts

Global Context **Contexts** Context Pools Settings

Label	Description	Tags	Operations
Bonnici Key Facts & Value Propositions	Business features, benefits, testimonials and target audiences	Drupal	Edit ▼
Landing Page Funnel Rules	Guidelines for creating landing pages tailored to three marketing funnel stages (top, middle, and bottom) with specific component structures	Drupal	Edit ▼
UI UX guidelines	General guidelines for our design system.	Drupal	Edit ▼

[HTTPS://WWW.BONNICI.CO.NZ/BLOG/DRUPAL-AI-NATIVE-PAGE-BUILDING-CANVAS-AI-CONTEXT](https://www.bonnici.co.nz/blog/drupal-ai-native-page-building-canvas-ai-context)

So we've talked about context a lot, so how would we as content strategists and content authors manage it? There are lots of different ways to approach this, but there is a project in the Drupal community that is really interesting called the Context Control Center. What it does is it allow you to define different kinds of contexts and select what situations they apply in. Then anything that hooks into Drupal's AI capabilities will uses these contexts as necessary. I just want to acknowledge that all the screenshots you're about to see here come from a blog post by an agency in New Zealand called Bonnici which I will link to at the end. I wanted to demo this myself but I had some issues getting it all working. Like a lot in the AI space it is still highly experimental.

So as an example here, you can see they have defined some broad overview content about the company and its work, rules for creating landing pages, and general UI/UX/Content guidelines. This is exactly the kind of stuff the system needs. These are defined as "Context Entities" in Drupalspeak, but in reality they're just text files written in Markdown. (By the way if you all need a new skill to pickup you should learn Markdown. Its kind of dumb but its growing wildly as the language of AI for text files.)

```
## Business Overview

### Key Facts
- New Zealand-based Drupal development agency
- Enterprise experience across government, healthcare, education
- Open source commitment - we contribute back to Drupal
- Boutique agency with senior-only team

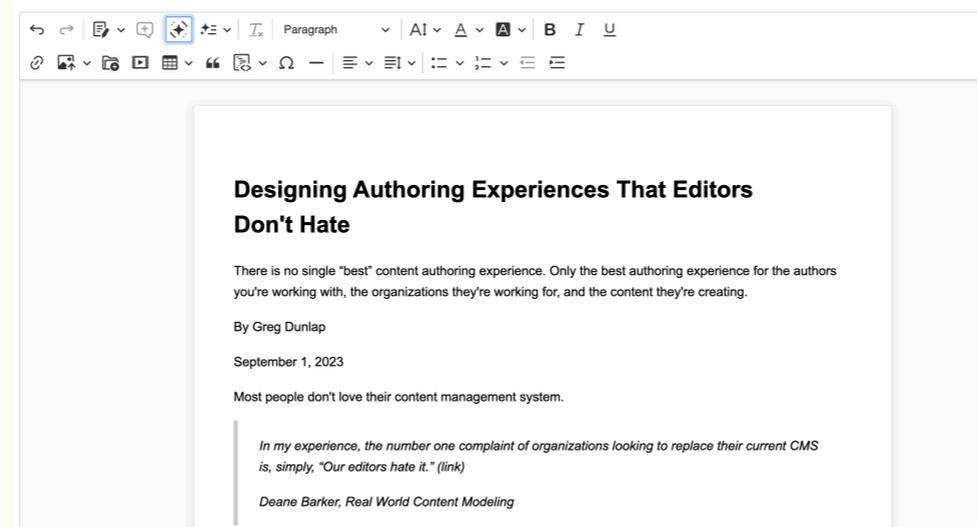
### Value Propositions
- Trusted expertise: 15+ years of Drupal development
- Enterprise-grade: We've built systems that serve millions
- Long-term partnership: Ongoing support, not just project delivery
- Transparent process: You see what we build, how we build it

### Testimonials & Social Proof
- Include client quotes where relevant
- Reference recognisable organisations when appropriate
- Let results speak - specific metrics over vague claims
```

So here's an example of one of these files. As you can see its a pretty simple description of some general company information. Who are we, how do we provide value? Very simple, but as you can see this is the kind of thing it would be great to have a content person write because you're trying to articulate information using plain language in a way that is easy to understand and organized well. I would probably want to see more detail than this personally, and then you can write contexts for things like your style guide, voice and tone rules, anything that guides how the system should work with content and UI. I feel like this is really cool, and I'm sure other CMS makers are working on similar stuff but this is a place where we can really provide a lot of value, and as I said, it all relates back to authoring because having this stuff defined means that any work the AI does on the content side is going to be better and more applicable.

So having said all that lets look at some ways that, once you have this context in place, AI can start to help content authors.

EDITOR TOOLING



A lot of the work in AI and authoring is happening within visual editors. I'm sure you have all encountered this kind of interface. It is generally attached to the body or article text field and offers various formatting options. For instance, this is a tool called CKEditor which is integrated into a lot of CMS systems. It has started building AI tools into its premium product, and this is something I can actually demo.

Go to demo
Show the review tools. Describe it as like having Grammarly built into your editor, however this is extensible, so a CMS can build additions to it for instance to ensure that its own contextual information gets brought into the fold. That is the real power behind it. It is not reviewing content based on some generic idea of what is "good" or "bad", it is specifically tied to your organization's rules and styles. Another thing you could work into these contexts is accessibility checking, not just generic but again in response to your organization's specific rules.

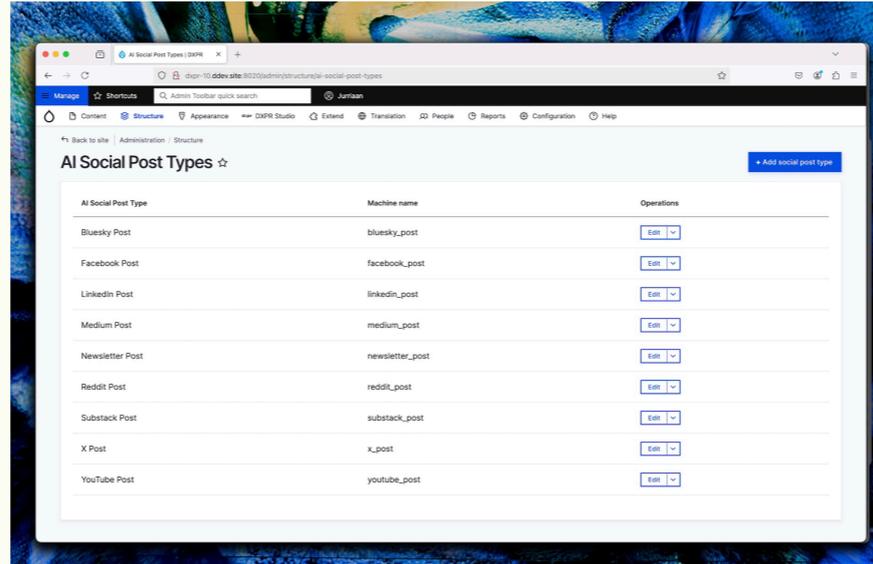
While we're talking about accessibility, there's been a lot of talk about auto-generating alt text but personally I find that to be a less compelling case. AI generally describes images very generically ("This is a white building on a hill") because it doesn't have the context of what they are ("This is the Iowa state capitol building"). Now, there's an argument that the generic alt text is better than nothing, but I'm just not very excited by any of that.

Note that these tools are limited to the context they live in - if there are other fields surrounding this body, or multiple CKEditor fields on a page, it won't necessarily know about them or integrate with them. This is really only useful for article types of content (which admittedly is quite a lot of content on the web)

Go into the chat and ask it to summarize as if generating a meta-description. Summarization is actually a kind of generation I can get behind. A lot of authors, particularly those who don't come from content backgrounds, have a terrible time summarizing content in a way that makes sense and is readable. They will often just copy the first paragraph in as a meta-description and call it a day. In these cases, where the context of an article can combine with the context of the system, AI has a lot of

information at hand to create something useful.

SUMMARIZATION



And then we can take that a step further in the example I'm showing here, a Drupal module that generates specific versions of posts that are optimized for different social networks. Not only will it write concise summaries tied to those systems' needs, but it will also recrop photos to specific aspect ratios and the like. This is the kind of drudgery I *want* AI helping me out with.

So these are just three examples I see where AI is helping to create a positive experience for authors and strategists that live inside the more traditional authoring interface. There are lots more, for instance there's a lot of interesting work going on with automatically applying taxonomies to content or organizing it in different ways. There's a lot of work happening in content governance doing things like continuous audits, identifying duplicate content in realtime, and not just broken link detection but broken content detection - referring to out of date products, ideas, past marketing campaigns, etc. I saw a demo of a new SAAS tool the other day that does intelligent regulatory compliance checking, preventing publishing of content that doesn't comply and keeping logs offsite to prevent them from being changed or deleted.

But to me, honestly, the more interesting work is happening *outside* the CMS, tools that link into it and interact with it. There is a technology that enables this called MCP which I'm not really going to get into, and getting this to work does require some configuration on both the CMS side and your side, but this work is really interesting and I'm going to demo it here.

OUTSIDE

Demo:

First go into Umbraco and show what it takes to create a simple event. (Create event, enter fields, enter location, optionally enter guest)

Then go into Claude code and do these actions

Create an event (pinball tournament)

Create a series of events (pinball tournaments for the next two months)

Create an event (paste from a URL - https://www.meetup.com/content-strategy-seattle/events/313175436/?eventOrigin=group_upcoming_events)

Import locations from a URL (migration use case - <https://www.va.gov/puget-sound-health-care/locations/>)

Word Document

These are straightforward examples, but I think there's so much potential here. For instance, in all these examples we have said "Create content of a type using (a description, a word doc, an external url, etc.) One of the things I spend a lot of time on in CMS implementations is describing content types so that authors will know which one to choose. As I often say, rarely will an author put a blog post into an event, but the reverse happens a lot. Or what is the difference between a blog post, an article, and a press release? These are subtle distinctions, and I often spend a lot of time describing those use cases within the CMS. What if we had a world where it was the other way around? You just drop your content in, and AI uses the context in the system to determine which content type is the correct one? This could be huge for authors and organizations alike. Less work for authors, more accurate content. I'm really excited to see where all this goes.

CONCLUSIONS

Conversational interface for content entry is legitimately the most impressive and exciting thing I've seen in the world of authoring tools since the introduction of the visual editor. I truly believe that it is going to make so many people's lives easier. Will it replace the CMS admin experience? Unlikely. Will it provide a pathway in for a great many people who would otherwise struggle? Absolutely. It is more intuitive, more accessible, and easier to understand. These tools are in their infancy but they are only going to get better. The amount of change and growth in just the last six months has been incredible. I understand that there is a lot to have about the way that AI is being pitched and implemented at organizations across the world. Lets face it, there is a lot more bad than good happening. However that doesn't mean these tools are not useful, and can't be brought to bear on problems authors have every day.

DESIGNING CONTENT AUTHORING EXPERIENCES

A book for the designers, strategists, and developers who build and maintain content management systems. With practical examples and best practices, the book will show you how to create content management systems that support authors, so that authors can better serve their audiences.

<https://authoringexperience.com>

Designing

CONTENT AUTHORING

E X P E R I E N C E S

By Greg Dunlap