

# GitHub Spec Kit: A Guide to Spec-Driven AI Development

By IntuitionLabs.ai • 10/7/2025 • 15 min read

spec-driven development

github spec kit

ai code generation

software specification

ai assisted development

github copilot

prompt engineering

developer tools



# Deep Dive into Spec Kit: Specification-Driven AI Development

Developers today increasingly use [AI assistants](#) (like GitHub Copilot and [ChatGPT](#)) to write code, but they often encounter **limited results**. As one GitHub blog explains, the common approach ("vibe-coding") is to sketch a goal and let the AI spit out code that *"looks right, but doesn't quite work."* ( [github.blog](#)). In practice, AI-generated code often fails to compile or misses the real intent. This isn't due to weak AI models, but our workflow: we've been treating AI like a search engine, rather than a careful partner ( [github.blog](#)). The models "excel at pattern recognition but still need unambiguous instructions" ( [github.blog](#)). In short, **missing context** is the problem: vague prompts yield unreliable code, and developers spend time debugging AI's guesses.

Spec-driven development (SDD) proposes a better approach. Instead of **writing code first and docs later**, we put a formal **specification up front** ( [github.blog](#)). That means defining *what* the software must do – user stories, acceptance criteria, flows, etc. – before writing code. The specification becomes a *living contract* and " [shared source of truth](#)" ( [github.blog](#)). In SDD, you capture the "why" behind every feature in version-controlled form, rather than letting design decisions hide in code or siloed documents ( [developer.microsoft.com](#)) ( [developer.microsoft.com](#)). For example, without a clear spec, a backend engineer, frontend engineer, and product manager might each assume a different behavior for a "notifications" feature – leading to messy rework ( [developer.microsoft.com](#)). Spec-driven development avoids that by surfacing all assumptions early. The spec **evolves with the project** as new insights emerge, so updating requirements becomes as routine as refactoring code ( [developer.microsoft.com](#)).

GitHub's **Spec Kit** is a new open-source toolkit designed to bring SDD practices to life using AI ( [github.blog](#)). Rather than tacking on documentation at the end, Spec Kit makes the spec the *center* of development ( [github.blog](#)). It provides a structured workflow and templates so that your AI assistant ("coding agent") always codes to the spec. As GitHub puts it, teams using Spec Kit "outline the concrete project requirements, motivations, and technical aspects" *before* handing tasks to the AI ( [developer.microsoft.com](#)). In other words, you feed the AI a detailed spec (and plan) and let it *build exactly what was needed* ( [developer.microsoft.com](#)). This flips the script: the specification is created first and continuously refined by humans, and the AI fills in the implementation to match that spec.

## The Four Phases of Spec-Driven Development

Spec Kit organizes work into **four clear phases** ( [github.blog](#)). You don't move to the next phase until the current artifact is reviewed and accepted. Each phase is a handoff between the human



(driving intent) and the AI (doing the writing):

1. **Specify:** The developer writes a high-level description of *what* needs to be built and *why* (goals, users, problem context). The AI then generates a **detailed specification** document. This spec covers user stories, acceptance criteria, workflows, and success metrics – *not* low-level code. It focuses on user experience and outcomes. For example, who is the user? What problem are we solving? How will they interact with the feature? The AI fleshes out all these details. ( [github.blog](#)). Crucially, this spec is a living artifact: you refine it as you learn more about users and constraints.
2. **Plan:** Next, the developer specifies **technical constraints** (desired tech stack, architecture style, compliance needs, performance targets, etc.). The AI generates a comprehensive **technical plan** based on these inputs. If your organization uses React and Node.js, that goes in here. If you need OAuth, [HIPAA compliance](#), or specific third-party systems, those details go here. The AI can even produce alternative plans for comparison (e.g. one plan using microservices vs. another using serverless). The key is that the AI *knows your rules and context* before writing any code ( [github.blog](#)).
3. **Tasks:** With a finalized spec and plan, the AI breaks the work into **concrete tasks**. These tasks are small, reviewable units (e.g. “Create a user registration endpoint that validates email format” rather than “build authentication”). Each task solves a specific piece of the specification. Breaking work down like this is akin to test-driven development for the AI: each task is something the AI can complete and validate independently ( [github.blog](#)). This keeps the AI on track and makes reviews manageable.
4. **Implement:** Finally, the AI assistant implements each task one by one (or in parallel). For each task, the developer reviews the generated code. Because the spec and plan guided the AI, it “knows what it’s supposed to build because the specification told it... and it knows exactly what to work on because the task told it” ( [github.blog](#)). Instead of inspecting huge AI-generated diffs, you see precise changes tied to clear requirements. Every code suggestion is traceable back to the spec, making validation faster.

This four-phase process ensures that **human insight drives AI coding**, not the other way around. At each step, you critique and refine the AI’s output: Does the spec capture what the user truly needs? Does the plan fit business constraints? Are tasks clear and complete? These built-in checkpoints catch misalignments early, long before they turn into bugs or rewrites ( [github.blog](#)).

## Using the Spec Kit Toolkit

Spec Kit provides a command-line interface (CLI) called `specify` to set up your project for SDD ( [developer.microsoft.com](#)). For example, you can run:

```
uvx --from git+https://github.com/github/spec-kit.git specify init <PROJECT_NAME>
```

This single command bootstraps the SDD scaffolding. It creates two key directories in your repo:



- **.github** – Contains prompt templates for your chosen AI agent (e.g. GitHub Copilot). Files like `specify.prompt.md`, `plan.prompt.md`, and `tasks.prompt.md` tell the agent how to generate each artifact.
- **.specify** – Contains the core SDD templates and scripts. For instance, `spec-template.md`, `plan-template.md`, and `tasks-template.md` are the skeletons that the AI fills in. There's also a `constitution.md` (setting project principles) and utility scripts for your OS (PowerShell or Bash) ([developer.microsoft.com](https://developer.microsoft.com)) ([developer.microsoft.com](https://developer.microsoft.com)).

Because the CLI downloads from the Spec Kit GitHub repo, you can use the official templates out of the box or customize them. Spec Kit is cross-platform: it provides shell scripts for Unix-like systems and PowerShell scripts for Windows, so you can use it in any IDE or CI/CD pipeline. The result is that Spec Kit **integrates into existing workflows** (“low friction” adoption) ([www.softwarezen.com](https://www.softwarezen.com)). You keep working in your repo; Spec Kit simply adds a framework to guide the AI through the SDD process.

## Benefits of Specification-First AI Development

Shifting to a specification-driven workflow unlocks several advantages:

- **Early Error Prevention:** By capturing requirements in detail upfront, teams prevent miscommunication. Without a spec, different developers might make conflicting assumptions; a clear spec aligns everyone (and the AI) from the start ([developer.microsoft.com](https://developer.microsoft.com)). For example, designers, frontend, backend, and PMs all derive feature behavior from the same spec, avoiding the classic “I thought you meant X” problem halfway through a sprint.
- **Living Documentation:** The specification becomes a living document that evolves with the project. It is version-controlled, reviewable, and always up-to-date. This means architectural decisions and business rules are explicitly recorded rather than buried in code. Developers can refactor code without losing context, because the spec tracks changes just like code commits ([developer.microsoft.com](https://developer.microsoft.com)).
- **Massive Efficiency Gains:** Investing time in precise planning pays off hugely. As one expert observed, doing deep requirements and architecture work up front can save an order of magnitude of effort later – “Every hour spent [on planning] saves 10 hours of rework” ([aroussi.com](https://aroussi.com)). Gone are the days of “Big Rewrite Monday” because of misunderstood requirements.
- **Role Shift for Developers:** The developer’s role changes from typist to architect/validator. Rather than writing boilerplate, you craft prompts, reviews, and designs. In effect, you become a *prompt engineer* and *quality controller*. One enthusiast put it: “**You’re now a prompt engineer, not a code writer. A quality controller, not a typist. An architectural guardian, not an implementer.**” ([aroussi.com](https://aroussi.com)). The AI handles routine implementation; developers focus on high-level thinking and ensuring fidelity to intent.

- **Flexible Variants:** Because the spec is decoupled from code, it's easy to generate multiple implementations or iterate designs. For instance, you could ask the AI to produce two versions of a component (say, one in Rust and one in Go) from the same spec to compare performance ( [developer.microsoft.com](https://developer.microsoft.com)). Or experiment with different UI mocks: change the Figma link in the spec and have the AI regenerate the corresponding pages. This kind of experimentation would be tedious if done by hand, but becomes straightforward when the spec is the single input to the AI.
- **Quality Improvement:** Overall code quality tends to improve. When the AI has clear requirements and smaller tasks, it produces more reliable code. The AI knows exactly *which part of the spec* it's implementing, reducing tangential mistakes. As a result, teams see fewer unexpected surprises and cleaner code that "does what we asked" ( [github.blog](https://github.blog)).

In short, specification-driven AI development turns software engineering into a feedback-driven process. You continuously steer the AI with explicit docs, rather than playing catch-up later.

## Spec Kit and the AI Ecosystem

Spec Kit is agent-agnostic by design. It works with any modern coding AI. Currently provided templates support GitHub Copilot, Google's Gemini (via CLI), Anthropic's Claude Code, and others ( [github.blog](https://github.blog)). Because the Spec Kit CLI just scaffolds files and prompts, you can plug it into existing projects without locking into a proprietary system. For example, a team using VS Code with Copilot can simply run `specify init`, and start using the Spec Kit prompts; in effect Copilot becomes a spec-driven assistant. Likewise, teams on JetBrains or CLI-based environments can integrate the templates.

Industry guides note that adopting Spec Kit involves minimal overhead: it "slots into familiar IDEs" and workflows ( [www.softwareseni.com](https://www.softwareseni.com)). You don't need a new environment or learning a new tool: it's mainly a mindset shift plus a ready-made framework. As AI coding becomes ubiquitous, toolchains around it are emerging; Spec Kit is one of the first comprehensive attempts to standardize how we do it.

## The Future: AI-Assisted Development, Specification by Default

Looking ahead, many experts agree that **specification-first development is the future of AI-assisted coding**. Several technical and industry trends point this way:

- **Massive Context Windows:** Modern language models can now handle extremely long inputs (200K+ tokens ( [www.softwareseni.com](https://www.softwareseni.com))). This means they can realistically consume entire requirement documents or API specs all at once. In fact, researchers note that LLMs "*understand formal specification formats like OpenAPI, JSON Schema, and structured documentation*" ( [www.softwareseni.com](https://www.softwareseni.com)). This opens the door to feeding stable, machine-readable specs (not just chat prompts) into the AI.

- Standard Specifications as Artifacts:** We are likely to see traditional spec formats (OpenAPI for APIs, JSON Schema for data, user-story markups, etc.) become first-class citizens in the development pipeline. Industry voices predict that *"specifications are becoming standard development artifacts"* and that tools will emerge around them ( [www.softwareseni.com](http://www.softwareseni.com)). For example, your CI/CD might automatically regenerate and deploy code whenever a spec file changes, treating the spec as source of truth. Documentation, tests, and even infrastructure can be generated from these formal specs, reducing duplication.
- AI in the Loop for Specs:** AI won't just consume specs; it will help write them. Future tools may automatically suggest missing requirements, generate edge-case tests, or refine ambiguous stories. But the key insight is that **high-quality specs unlock AI's potential**. As one author put it: *"AI can generate high-quality code, but only if we give it high-quality specifications"* ( [foojay.io](http://foojay.io)). In practice, this means human teams will invest more effort up front in requirements engineering and less in manual coding. The developer becomes a designer of requirements and constraints, and the AI acts as a fast implementer.
- Focus on Architecture & Validation:** As repetitive coding falls to AI, human developers will focus on what AI can't easily do: deep architecture, system design, security boundary setting, and user experience. The SoftwareSeni guide predicts a world where *"human developers will focus on architecture, requirements, and validation rather than manual implementation."* ( [www.softwareseni.com](http://www.softwareseni.com)). In other words, developers curate and validate the AI's output rather than writing every line themselves.
- Reduced Rework and Faster Delivery:** Companies that embrace specs and AI could build software much faster and with fewer bugs. As one enthusiast notes, this shift is not about *replacing developers*, but *"using AI to eliminate the tedious, error-prone work so we can focus on what matters: understanding business needs and designing systems"* ( [foojay.io](http://foojay.io)). The future code pipeline might routinely consist of: **1)** write a detailed spec, **2)** trigger AI code generation, **3)** run tests and reviews, **4)** update spec as needed and repeat.
- Emerging Standards & Ecosystems:** We may also see industry standards arise around *"specification as code."* Just as DevOps popularized *"Infrastructure as Code,"* we might get *"Specification as Code"* toolchains. Some early projects (like the *ai-development-specifications* repo ( [github.com](https://github.com))) already document complete SDD practices. Major platforms (GitHub, Amplitude, etc.) might build first-party support to store and version specs alongside code.

In short, the **trajectory is clear**: the next phase of AI-assisted development is specification-driven. By treating specs as living, executable artifacts, teams can harness AI's raw power while anchoring it to real human intent ( [github.blog](https://github.blog)) ( [foojay.io](http://foojay.io)). As one blog puts it, *"The code is the easy part. Getting the requirements right is where the real value lies."* ( [foojay.io](http://foojay.io))

## Conclusion

Spec Kit exemplifies a fundamental shift in how we collaborate with AI on software. It codifies a spec-first philosophy that promises higher-quality results and more predictable outcomes. Instead of vague chat prompts, we create a structured dialogue: human-to-spec, spec-to-AI, AI-



to-code, and back. This process leverages the strengths of both parties – clear human judgment up front, and AI's fast implementation – while minimizing their weaknesses.

As AI tools continue to evolve, specification-driven development may well become the industry norm. In the coming years, we can expect **specifications to be treated as code artifacts**, baked into our workflows, and automatically enforced by AI assistants ( [www.softwareseni.com](http://www.softwareseni.com)) ( [www.softwareseni.com](http://www.softwareseni.com)). For developers, this means moving away from typing code line-by-line and toward architecting solutions at a high level. With Spec Kit and similar approaches, we're already seeing how giving AI good instructions (in the form of specs) makes all the difference. The future of AI-assisted coding is bright – provided we remember that *how* we communicate with AI (through clear specifications) is just as important as the AI itself ( [github.blog](https://github.blog)) ( [foojay.io](http://foojay.io)).

**References:** Spec Kit documentation and blogs ( [github.blog](https://github.blog)) ( [github.blog](https://github.blog)) ( [developer.microsoft.com](https://developer.microsoft.com)) ( [developer.microsoft.com](https://developer.microsoft.com)) ( [developer.microsoft.com](https://developer.microsoft.com)) ( [developer.microsoft.com](https://developer.microsoft.com)) ( [aroussi.com](https://aroussi.com)) ( [foojay.io](http://foojay.io)) ( [www.softwareseni.com](http://www.softwareseni.com)) ( [www.softwareseni.com](http://www.softwareseni.com)). These sources outline the spec-driven process, its benefits, and the emerging vision of specification-centric AI development.

---

## IntuitionLabs - Industry Leadership & Services

**North America's #1 AI Software Development Firm for Pharmaceutical & Biotech:** IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

**Elite Client Portfolio:** Trusted by NASDAQ-listed pharmaceutical companies including Scilex Holding Company (SCLX) and leading CROs across North America.

**Regulatory Excellence:** Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

**Founder Excellence:** Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

**Custom AI Software Development:** Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

**Private AI Infrastructure:** Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

**Document Processing Systems:** Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

**Custom CRM Development:** Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

**AI Chatbot Development:** Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

**Custom ERP Development:** Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

**Big Data & Analytics:** Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

**Dashboard & Visualization:** Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

**AI Consulting & Training:** Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at <https://intuitionlabs.ai/contact> for a consultation.



---

## DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will [IntuitionLabs.ai](https://IntuitionLabs.ai) or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

[IntuitionLabs.ai](https://IntuitionLabs.ai) is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by [Adrien Laurent](#), a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.

© 2025 [IntuitionLabs.ai](https://IntuitionLabs.ai). All rights reserved.