PixiTale

AI-powered personalised children’s story books

Serhat Keskin

June 2026

Contents

1 The Problem

Children aged 2–8 spend growing hours in front of screens, yet the vast majority of that time is passive — consuming algorithmic, ad-driven content that does nothing to spark imagination or reinforce family values.

Traditional picture books are beloved but static. A story about *someone else's* character, in *someone else's* world, rarely holds a toddler's attention the way a story about *them* would.

Parents are left with a painful choice: hand over the tablet and accept low-quality content, or fight the screen-time battle with no winning alternative.

> **Core tension:** Children want the magic of digital media. Parents want the depth and safety of books. Nothing on the market delivers both.

2 The Solution — PixiTale

Figure 1 A generated story character — the child becomes the protagonist of their own adventure

PixiTale is a SaaS platform that turns a child's photo (and optionally a favourite toy's photo) into the hero of a fully illustrated, narrated picture book — generated end-to-end by AI in minutes.

Parents choose the story's language, age group, and pedagogical theme. PixiTale then:

1. Analyses the uploaded photos with **Vision AI** to build a consistent character description.
2. Generates an original **story script** (title, scenes, dialogue) via a large language model.
3. Renders **custom illustrations** for every page that keep the child's likeness scene-to-scene.
4. Produces **narration audio** with TTS, synced to karaoke-style highlighted text.
5. Assembles everything into an interactive **page-flip book** inside the app.

The result is a unique book the child can watch, listen to, and read — starring themselves.

3 Architecture

The platform is built on a **fully asynchronous, event-driven architecture** designed to handle heavy AI workloads without blocking the user experience.

3.1 Backend

- **Django 5.2 + DRF** — REST API, authentication, business logic, credit/subscription management.
- **Celery 5.4 + Redis 7** — distributed task queue. Character analysis, illustration generation, and TTS all run as parallel Celery tasks so a 6-page book's assets are produced concurrently rather than sequentially.
- **PostgreSQL 14** — relational store for users, books, pages, generation process state machines.
- **AWS S3** — object storage for all generated images and audio files.
- **Flower** — real-time Celery task monitoring dashboard.

3.2 AI Pipeline

All AI calls go through `backend/ai_services/`, an abstraction layer that makes providers swappable without touching business logic.

- **Google Gemini 2.5 Flash** — story text generation and image-prompt engineering.
- **Google Gemini 2.5 Flash (Vision)** — photo analysis to extract character features for consistency.
- **Gemini 2.5 Flash Image** — scene illustration with character reference injection.
- **Gemini 2.5 Flash Preview TTS** — natural-voice narration per scene.

3.3 Frontend

- **React 19 + Vite 6 + TypeScript** — fast SPA with hot module reloading in development.
- **Ionic 8 / Capacitor 7** — single codebase compiled to iOS, Android, and web PWA.
- **Tailwind CSS 3.4** — utility-first styling, responsive across device sizes.
- **react-pageflip** — realistic page-turn animation for the in-app book reader.

3.4 Infrastructure

Docker Compose orchestrates seven services locally (`st-backend`, `st-frontend`, `st-db`, `st-redis`, `st-celery-worker`, `st-celery-beat`, `st-flower`) behind an Nginx reverse proxy with TLS. The same compose topology maps directly to a production-ready container deployment.

4 App Walkthrough

4.1 Story Creation Wizard

The four-step wizard guides parents through character selection, optional companion item, story settings, and a free-text prompt. Completed in under a minute.

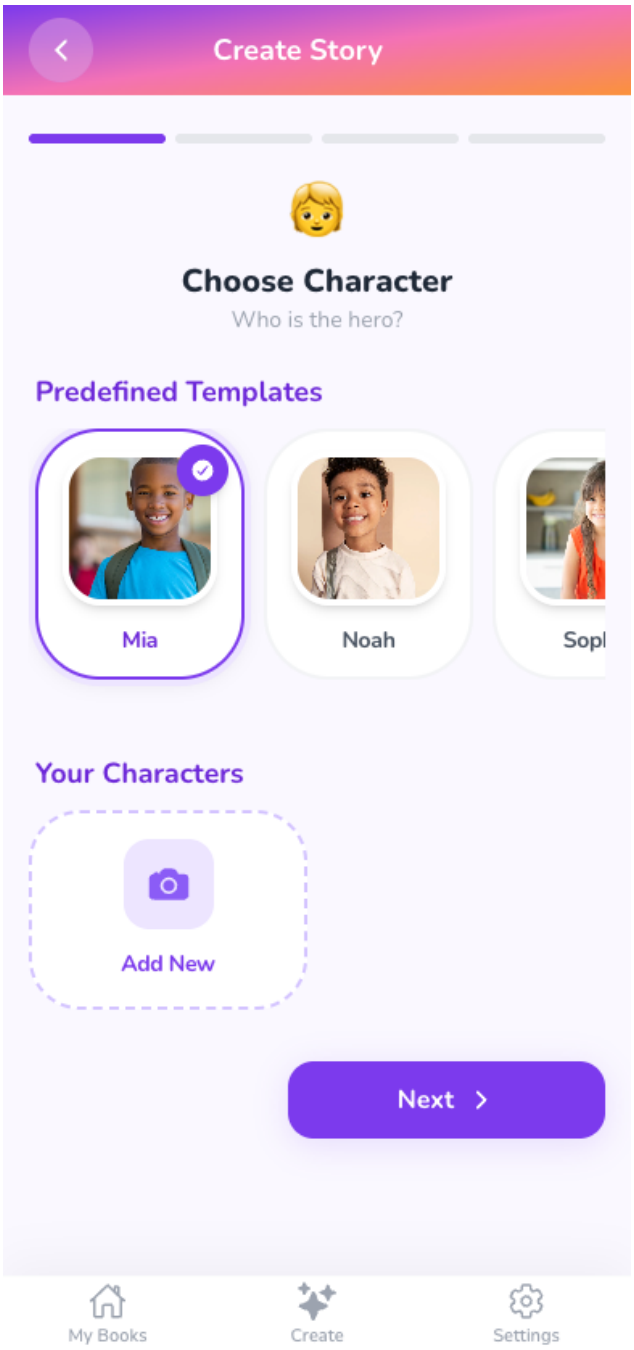

Figure 2 Step 1: Choose Character — predefined templates or custom upload

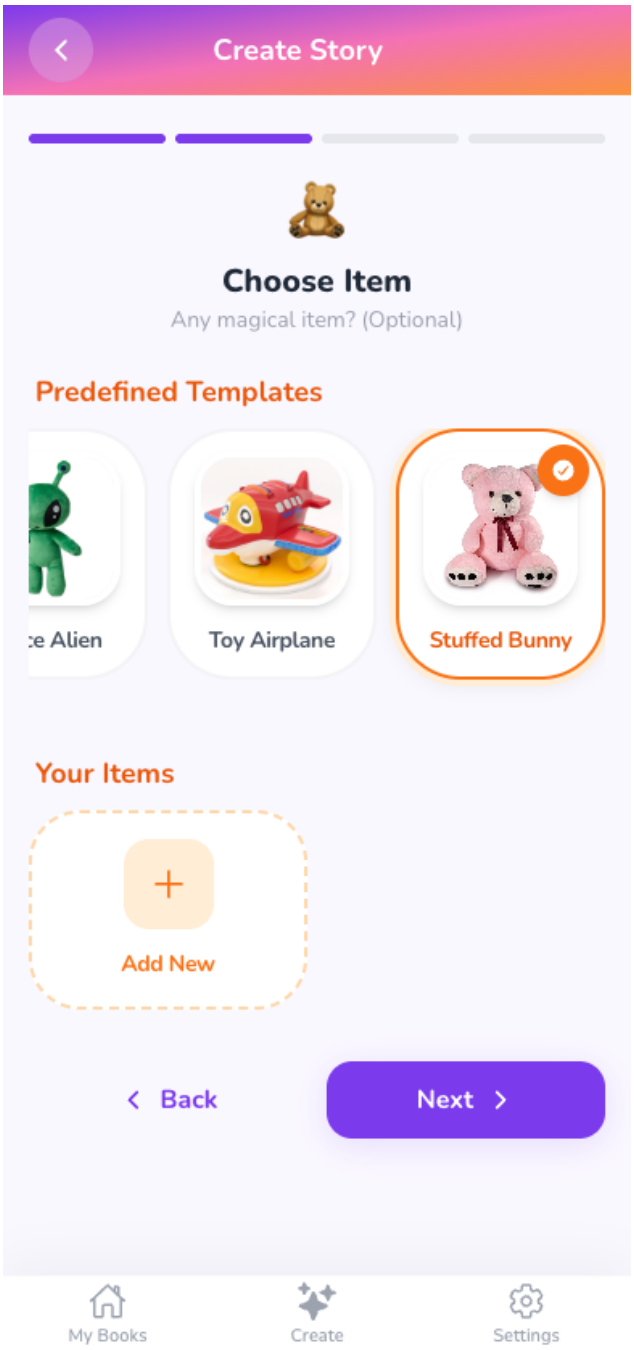

Figure 3 Step 2: Optional companion item (Stuffed Bunny selected)

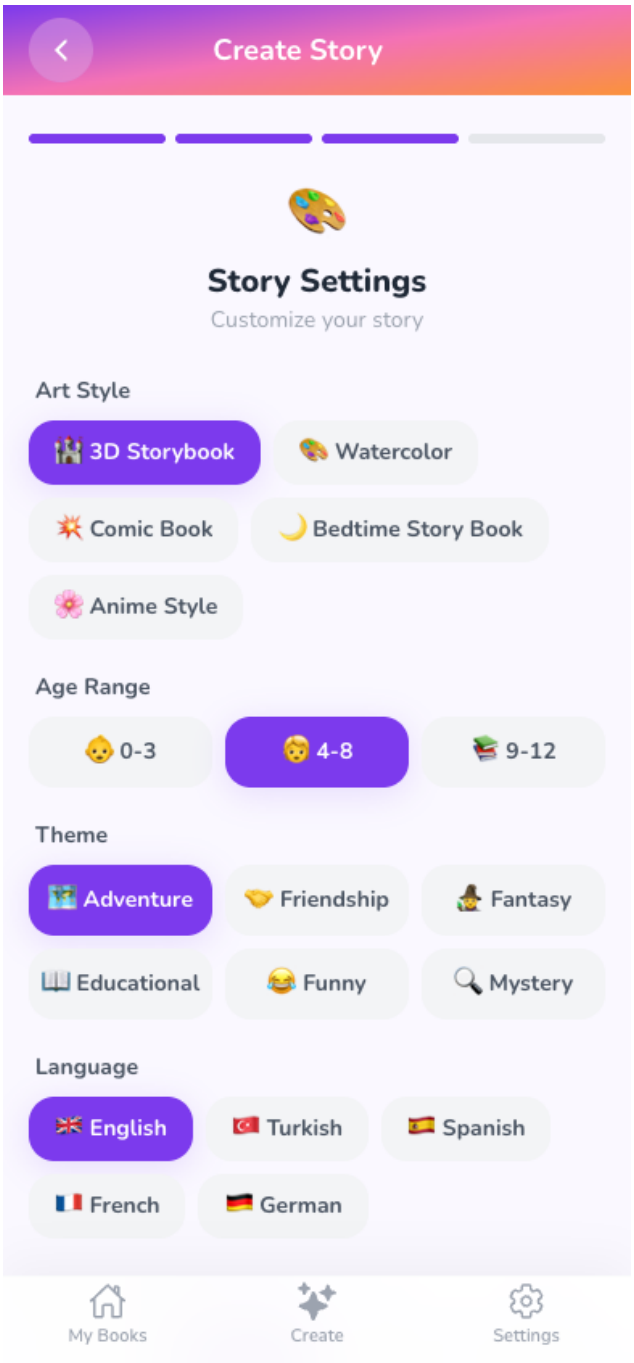

Figure 4 Step 3: Story settings — art style, age range, theme, language, length

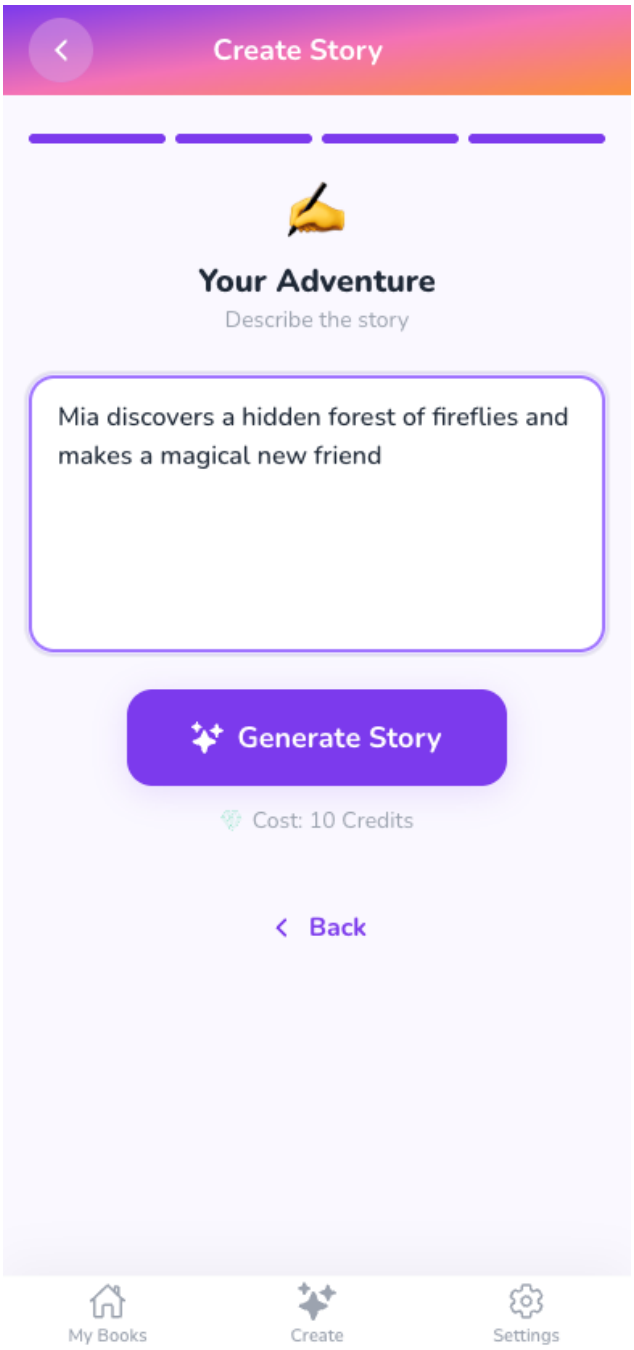

Figure 5 Step 4: Free-text adventure prompt and Generate Story button

4.2 AI Generation Pipeline

After submission, an overlay tracks the asynchronous pipeline in real time. Vision, LLM, image generation, and TTS tasks execute in parallel via Celery. A 5-page book with real AI illustrations completes in 60–90 seconds.

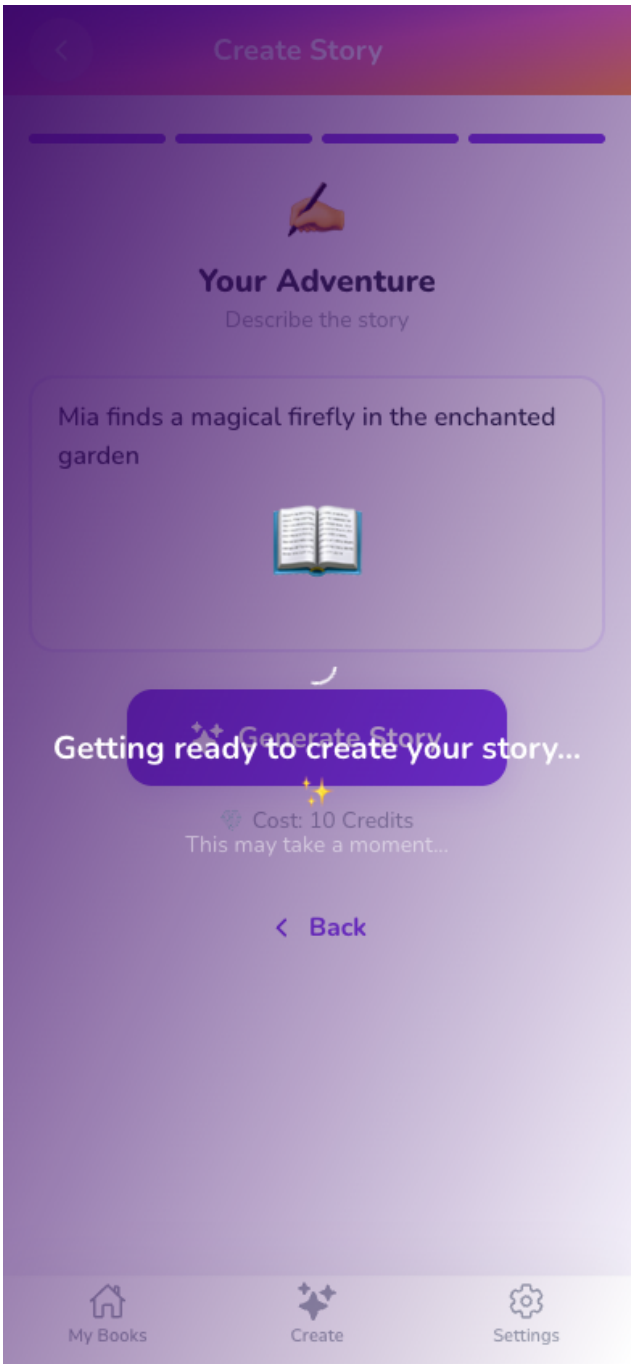

Figure 6 Real-time generation progress — “Getting ready to create your story...”

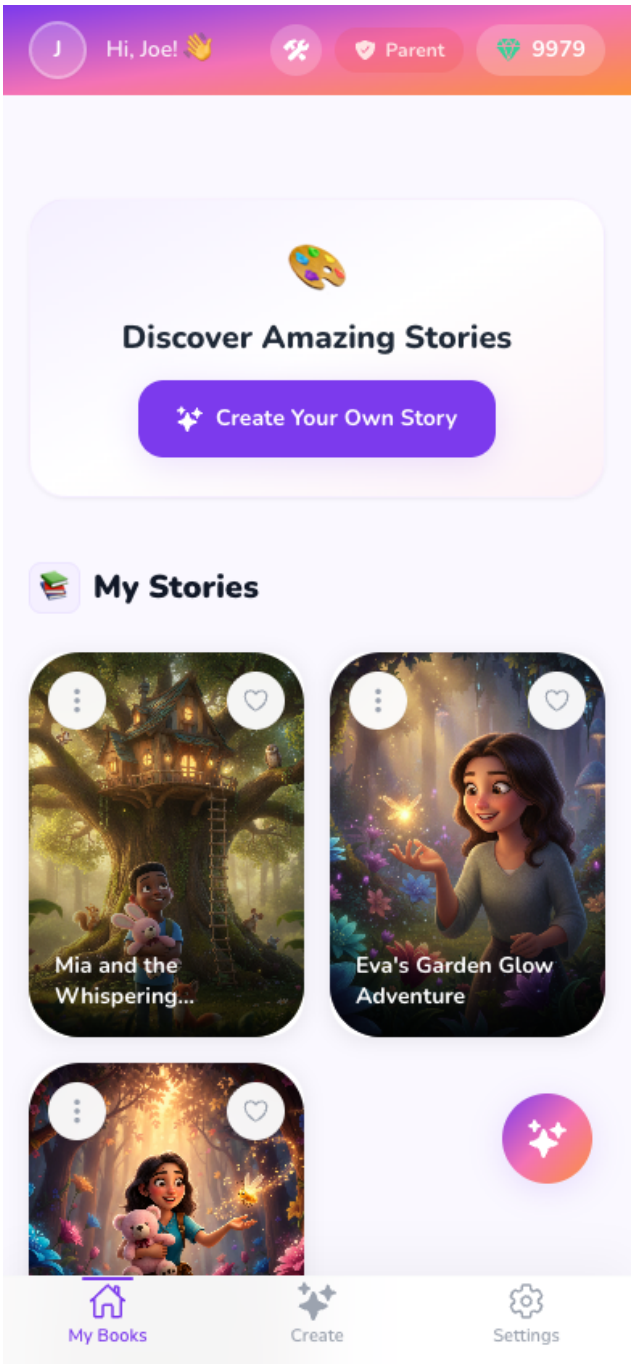

Figure 7 Home shelf after generation: *Mia's Firefly Adventure* with AI-generated cover

4.3 Page-Flip Book Reader

Completed books open in a page-flip reader with reading-mode selection, karaoke-style narration sync, and audio controls. Every illustration is uniquely generated by AI.

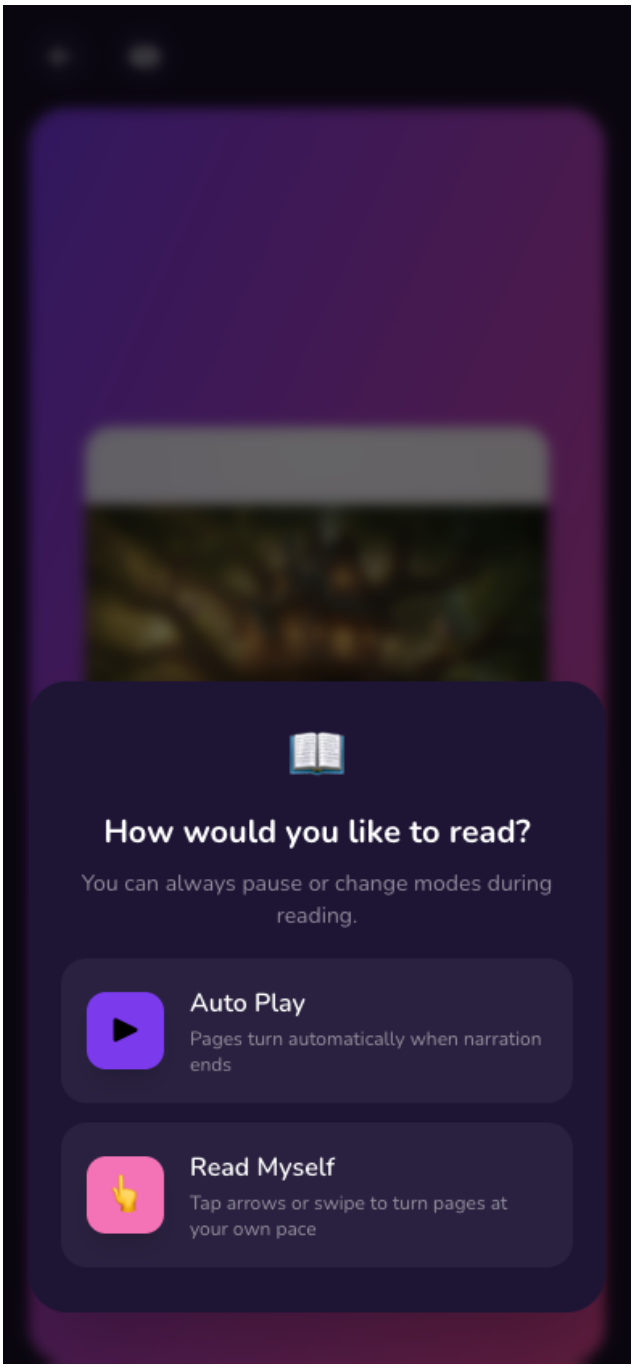

Figure 8 Reading mode selection: Auto Play with narration or Read Myself

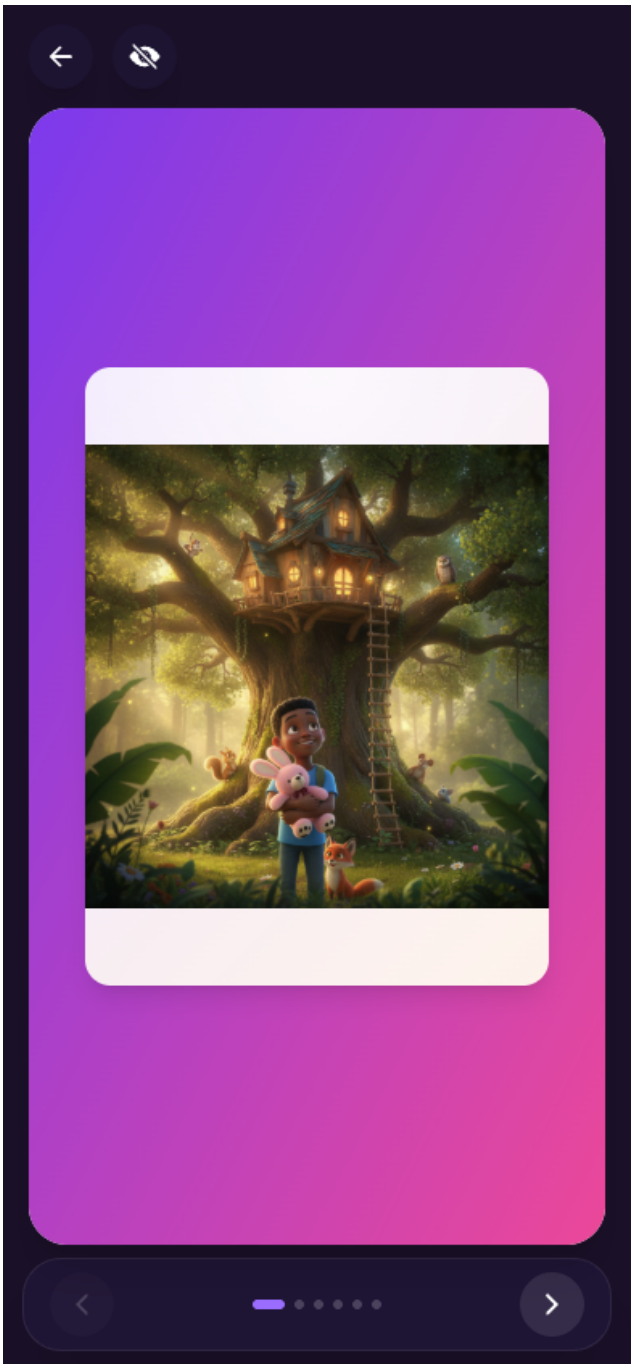

Figure 9 Book cover — AI illustration of Mia at the enchanted forest treehouse

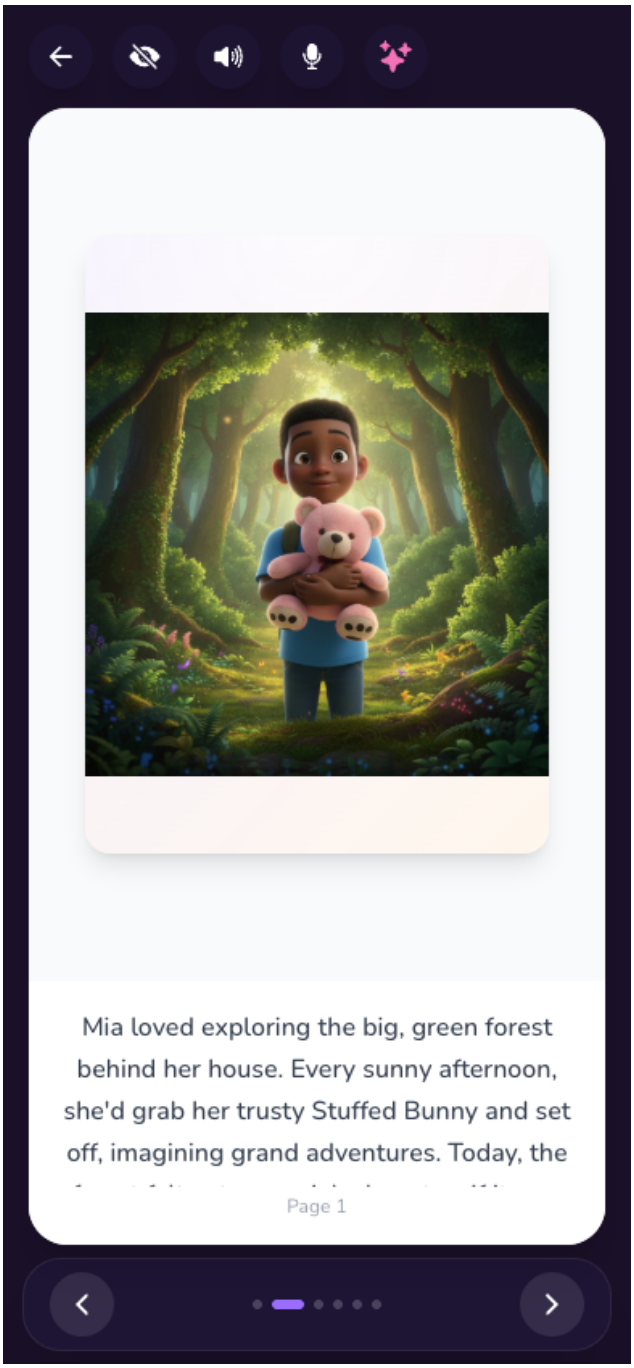

Figure 10 Page 1 — Mia and her Stuffed Bunny set off into the magical forest

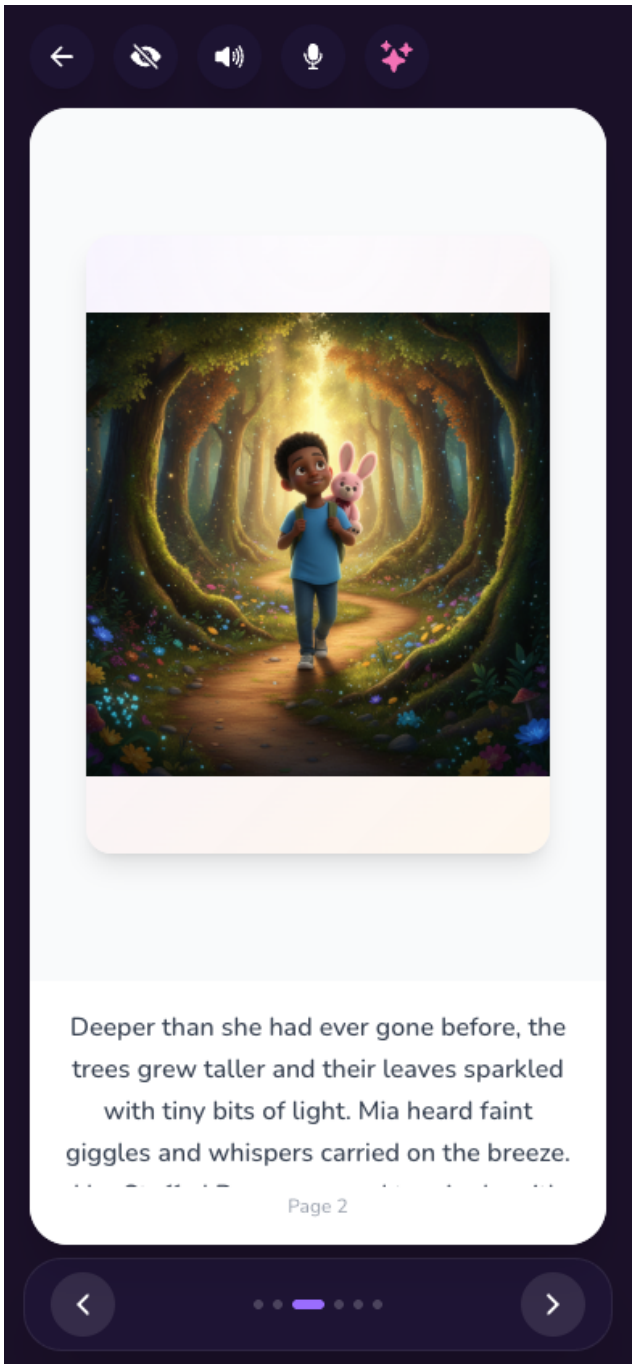

Figure 11 Page 2 — consistent character likeness maintained across every scene

4.4 Settings and Parental Controls

The Settings screen exposes credits and billing, language preference, and parental controls. Child Mode switches to an ad-free locked interface; re-entry requires a 4-digit PIN.

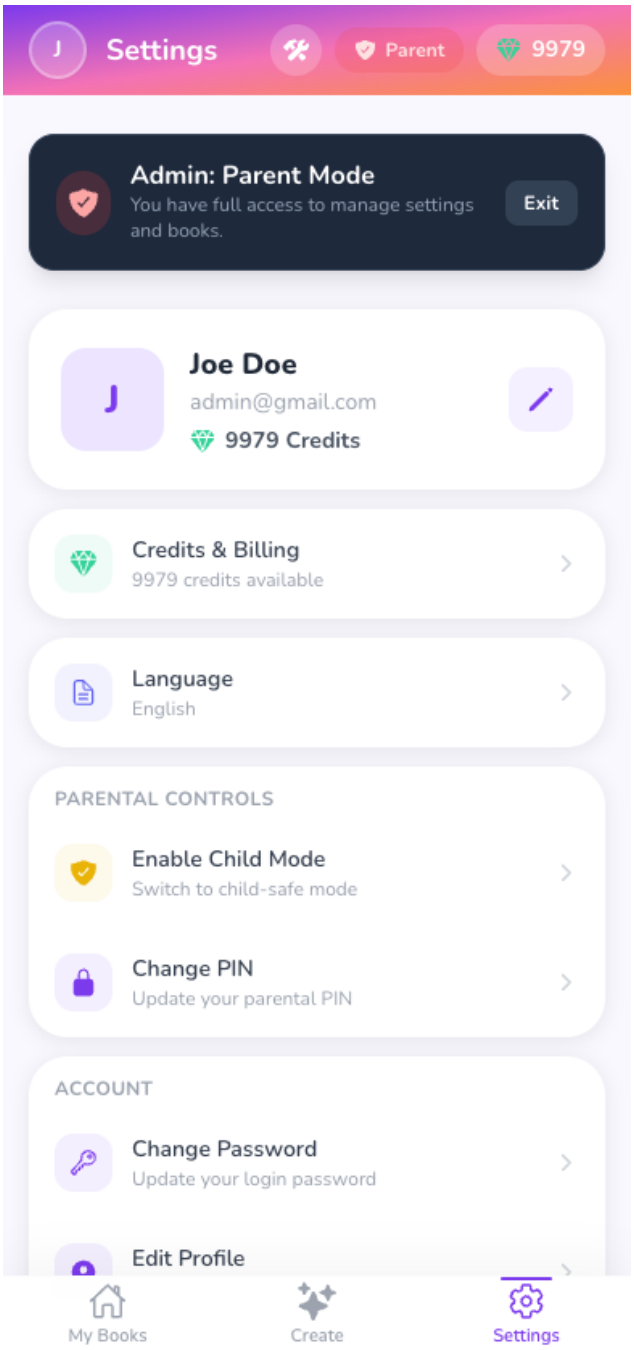

Figure 12 Settings page — profile, credits, language, and parental controls

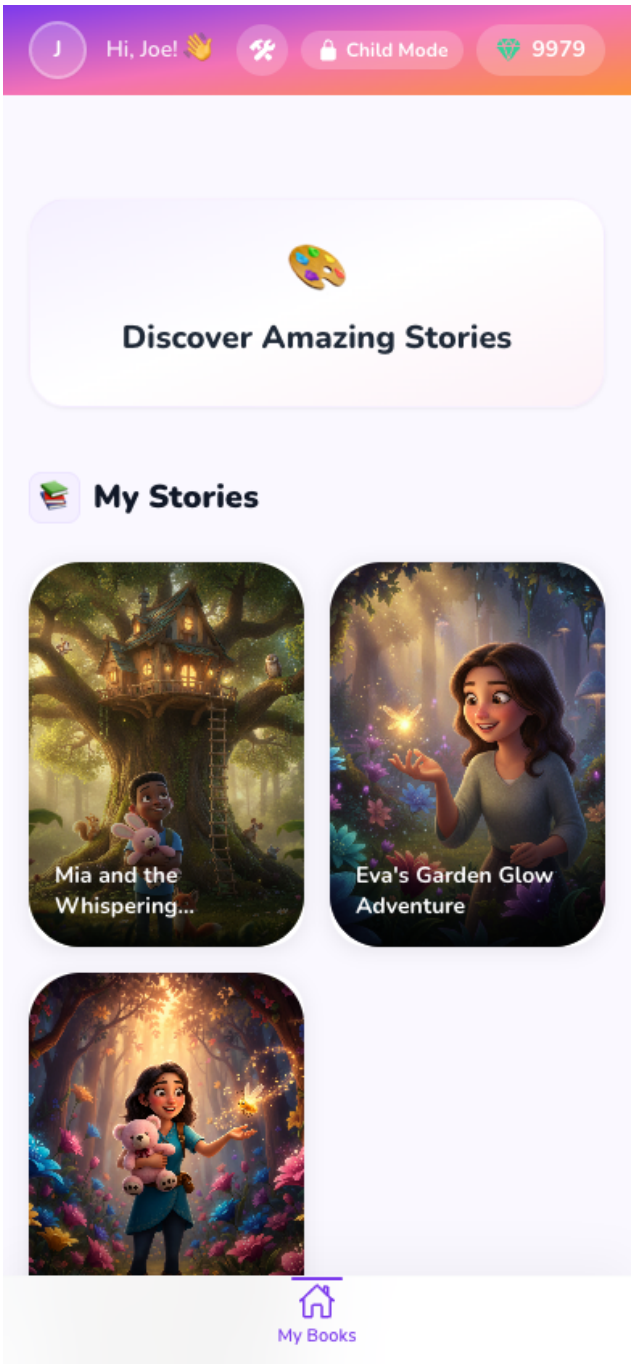

Figure 13 Child Mode home — *Mia's Firefly Adventure* in the ad-free book shelf

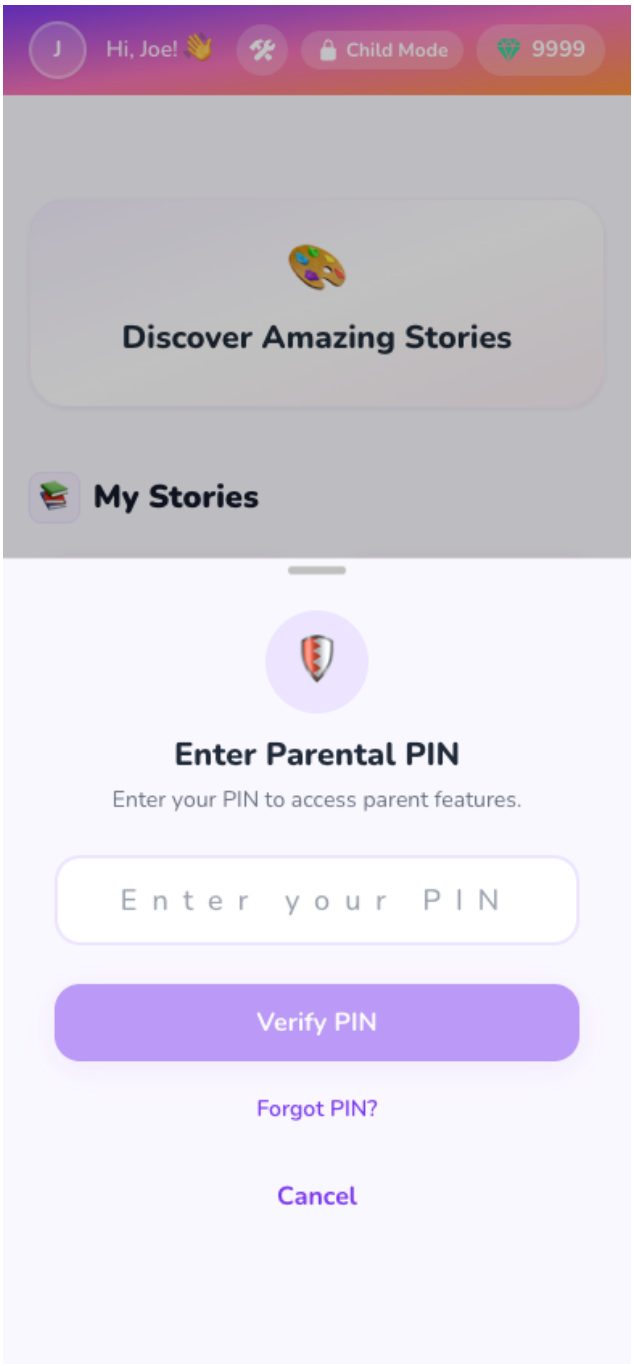

Figure 14 Parental PIN gate to exit Child Mode

5 Key Features

5.1 Visual Character Consistency

PixiTale's core technical challenge is keeping a child's face and clothing consistent across 6–10 generated illustrations. The platform achieves this through structured character descriptions extracted by Vision AI, fed as reference context into every image-generation prompt — delivering scene-to-scene coherence that competing solutions cannot match with naive prompt engineering alone.

5.2 Safe Kids Mode

A locked, ad-free child interface shows only approved books. A **Parental Gate** (PIN or biometric) controls all settings and purchases. No third-party ads, no external links, no algorithmic recommendations.

5.3 Multilingual & Multi-age

Stories are generated in the parent's chosen language (Turkish, English, German, Spanish, French supported at launch). Age group selection adjusts vocabulary complexity, sentence length, and illustration style density — appropriate from toddler (2–4) to early reader (6–8).

5.4 Flexible Business Model

- **Pay-as-you-go credits** — parents preview the text story for free; spend 1 credit to unlock illustrations and audio.
- **Monthly / annual subscription** — bundled credit allowance plus premium features (voice cloning, character library).

5.5 Privacy & GDPR / KVKK Compliance

Child photos are stored encrypted at rest, never shared with ad networks or used for AI model training. A one-click *Right to Erasure* flow permanently deletes all user data within 24 hours.

6 Tech Stack at a Glance

Layer	Technology
API / Backend	Django 5.2, Django REST Framework, Python 3.12
Task Queue	Celery 5.4, Redis 7, Flower monitoring
Database	PostgreSQL 14
Object Storage	AWS S3 (django-storages + boto3)
AI — LLM	Google Gemini 2.5 Flash
AI — Vision	Google Gemini 2.5 Flash (multimodal)
AI — Image Gen	Gemini 2.5 Flash Image
AI — TTS	Gemini 2.5 Flash Preview TTS
Frontend	React 19, Vite 6, TypeScript, Tailwind CSS 3.4
Mobile	Ionic 8, Capacitor 7 (iOS + Android)
Infrastructure	Docker, Docker Compose, Nginx, TLS

7 Project Links

- **Website:** pixitale.app
- **Developer:** Serhat Keskin