

SubtitleSync

AI subtitles in 30 languages + dubbing in 12 — cheaper than Rev, faster than Captions. Cost of multi-language localisation is the creator's bottleneck for international audience growth. \$19/month.

| | |
|---------------------|--|
| Category | Set 7 · Verticals & Creator |
| Customer | YouTubers + indie filmmakers + course creators + podcast creators wanting multi-language subtitles + dubbing |
| Monetisation | \$19/mo Solo (60 min/mo) · \$49/mo Pro (300 min/mo + dubbing) · \$149/mo Studio (1,000 min/mo + advanced) |
| Build effort | Med |
| Plan version | v1.0 — 2026-05 |

Executive Summary

SubtitleSync addresses the multi-language localisation cost-bottleneck for video creators. The structural opportunity: serious creators (YouTubers + indie filmmakers + course creators) have realised that international audience growth requires multi-language subtitles + dubbing — a 10-minute video subtitled in Spanish + Hindi + Portuguese + Arabic + Russian potentially reaches 3-5x larger audience. But cost: Rev.com human subtitles run \$1.50/min per language; human dubbing \$5-15/min. For a creator publishing 4 videos/week × 10 min × 5 languages, this is \$1,200-3,000/week — prohibitive.

Product: AI subtitles in 30 languages + AI dubbing in 12 languages. SubtitleSync uses combination of Whisper-class transcription + LLM translation tuned for video-content + ElevenLabs-class voice synthesis for dubbing.

Year-1 target: 3,500 paying subscribers generating █3.4 crore annual revenue against █62 lakh costs. Cash-positive month 3.

The Problem

Video creators wanting international audience growth need multi-language subtitles + dubbing. The economic gap: human-quality multi-language localisation (Rev.com subtitles at \$1.50/min + Voiceover artists at \$5-15/min) is prohibitive for typical creator economics. Result: most creators publish English-only + leave 3-5x audience growth on the table.

Existing AI options. YouTube auto-translated subtitles: free but low quality. Captions.app + Rask.ai + HeyGen + similar: exist but with various trade-offs (Captions weak on multi-language; Rask focused on translation but not native subtitle workflow; HeyGen focuses on avatar+dubbing with high-end pricing).

Market gap: focused multi-language subtitle + dubbing at moderate creator-friendly pricing with quality good enough for serious creator use.

The Solution

SubtitleSync's flow. Upload video (or paste YouTube URL); select target languages from 30 options. AI processes within 30-90 minutes: transcribes original language + translates to target languages + generates synchronised subtitle files (SRT + VTT) + optionally generates dubbed audio in 12 languages.

Quality optimisation. Translation: video-content-aware translation (preserves creator's tone + handles colloquial language + handles technical terminology with optional glossary). Subtitle timing: respects line-length conventions + reading-speed limits per language.

Dubbing (12 high-volume languages): high-quality voice synthesis with voice-style matching + lip-sync awareness where applicable.

Pricing tiers. Solo \$19/mo (60 min/month subtitles only). Pro \$49/mo (300 min/month + dubbing 60 min/month). Studio \$149/mo (1,000 min/month + advanced editing + team workspace).

Three structural differences from competitors. First, integrated subtitle + dubbing (vs. piecemeal). Second, creator-tier pricing (vs. enterprise pricing of HeyGen). Third, 30-language subtitle + 12-language dubbing breadth.

Market Opportunity

Global creator market with meaningful video output: ~15M creators. Subset doing multi-language localisation: ~500k currently; expanding to 2-3M as AI tools make it affordable.

At \$300/yr blended ARPU, SAM is \$150-900M growing rapidly. Realistic 4-year capture: 0.5-2% = \$4-18M ARR.

Adjacent expansion. Year 2: podcast localisation tier (similar workflow for podcast audio). Live-event interpretation (conference + webinar simultaneous translation). Year 3: enterprise corporate-training localisation tier.

Target Customer

Primary persona: a 31-year-old tech-tutorial YouTuber with 180k subscribers wanting to add Spanish + Portuguese + Hindi subtitles to grow international audience. Will pay \$19/mo Solo.

Secondary persona: a 38-year-old course creator with 12 published courses wanting dubbing in 4 languages for Udemy international expansion. Will pay \$49/mo Pro tier.

Tertiary persona: a 41-year-old indie filmmaker doing festival-circuit work needing professional-quality subtitles in 8 languages for international submissions. Will pay \$149/mo Studio tier.

Product

Upload: video file or YouTube/Vimeo URL.

Language selection: subtitles 30 languages; dubbing 12 languages.

AI transcription + translation + subtitle generation.

AI dubbing with voice-style matching.

Output: SRT + VTT + dubbed-audio files + optionally MP4 with burned-in subtitles.

Editor: in-app subtitle editing for refinement before export.

Glossary management: per-creator brand-terminology glossary for translation consistency.

Pro tier additions: dubbing + advanced editing + team workspace.

Studio tier additions: 1000 min/month volume + advanced editing tools + priority processing.

Technical Architecture

Frontend: Next.js + Tailwind.

Backend: Python on Hetzner cloud + GPU box for AI processing.

Transcription: Faster-Whisper Large-v3.

Translation: GPT-4o + Claude for video-content-aware translation.

Dubbing: ElevenLabs Multilingual + custom voice-style models.

Audio mastering: FFmpeg + Sox.

Payments: Stripe + Razorpay.

Business Model & Unit Economics

Three tiers. Solo \$19/mo (60 min subtitles). Pro \$49/mo (300 min subtitles + 60 min dubbing). Studio \$149/mo (1000 min subtitles + advanced + team workspace).

Conversion: 14-day trial converts at 22%. Distribution: 55% Solo, 35% Pro, 10% Studio.

Gross margin: 72%. Major costs: AI inference (subtitle ~\$0.04/min; dubbing ~\$0.40/min), infrastructure.

LTV: \$228 × 14 mo = \$319 (Solo); \$588 × 20 mo = \$1,176 (Pro); \$1,788 × 28 mo = \$5,006 (Studio).

Unit Economics (Year-1 base case)

| | |
|----------------------------------|-------------------------|
| Year-1 paying subscribers | 3,500 |
| Blended ARPU | \$118/mo |
| Year-1 revenue | \$410,000 (~₹3.4 crore) |
| Gross margin | 72% |
| CAC | \$60 |
| Year-1 all-in costs | ~₹62 lakh |
| Year-1 net contribution | ~₹2.1 crore |

Go-to-Market

Channel 1 — Creator-community (45%): YouTuber + indie filmmaker + course-creator communities + creator-tool ecosystem.

Channel 2 — Content + SEO (25%): comparison content vs. Captions + Rask + HeyGen.

Channel 3 — Creator-tool integrations (20%): partnerships with Riverside.fm + Descript + creator platforms.

Channel 4 — Paid acquisition (10%).

Roadmap (first 12 months)

- Month 1-3: MVP with subtitle generation in 10 languages + Solo tier. 300 subscribers.
- Month 4-5: 20 languages + Pro tier with dubbing in 6 languages, 900 subscribers.
- Month 6-8: 30 languages subtitles + 12 languages dubbing + Studio tier, 2,000 subscribers, ₹18 lakh MRR.
- Month 9-10: In-app editor + glossary management + team workspace, 2,900 subscribers.
- Month 11-12: 3,500 subscribers, ₹3.4 crore annualised.

Key Risks

- AI quality variance: subtitle + dubbing quality varies by language pair + content type. Mitigated by language-pair-specific tuning + clear quality expectation-setting.
- Captions / Rask / HeyGen competitive responses — possible. Mitigated by integrated subtitle + dubbing + breadth of language coverage + creator-friendly pricing.
- ElevenLabs dependency for dubbing: API cost + access changes. Mitigated by multi-provider abstraction + open-source TTS fallback (XTTS + similar).
- Voice-cloning ethical concerns: dubbing-in-creator's-voice raises consent + identity questions. Mitigated by clear consent workflow + voice-marketplace alternatives.

- Cultural / nuance translation issues: certain content doesn't translate cleanly. Mitigated by glossary management + manual-editing workflow + clear language-pair quality indicators.