

How many tokens Loom can save

Based on actual Codex and Claude use from September 10-17, 2026. The lower end of each range is the conservative estimate; the upper end is the realistic estimate.

16.6x-54.7x TOKENS SAVED FOR EACH TOKEN LOOM USES CONSERVATIVE TO REALISTIC	57.8%-82.5% REDUCTION IN TOTAL TOKEN USE CONSERVATIVE TO REALISTIC
---	--

Tokens saved for each token Loom used

Type of use	Tokens used by Loom	Conservative tokens saved	Realistic tokens saved	Conservative saved / Loom token	Realistic saved / Loom token
Codex interactive sessions	3,526,219	53,737,609	174,862,685	15.2x	49.6x
Codex automated sessions	2,309,852	42,631,411	141,848,769	18.5x	61.4x
Claude automated sessions	1,362,743	23,421,998	76,806,400	17.2x	56.4x
Combined	7,198,814	119,791,018	393,517,854	16.6x	54.7x

Reduction in total token use

Type of use	Tokens used with Loom	Estimated tokens without Loom conservative	Estimated tokens without Loom realistic	Conservative reduction	Realistic reduction
Codex interactive sessions	27,675,218	77,886,608	199,011,684	64.5%	86.1%
Codex automated sessions	33,018,064	73,339,623	172,556,981	55.0%	80.9%
Claude automated sessions	21,430,394	43,489,649	96,874,051	50.7%	77.9%
Combined	82,123,676	194,715,880	468,442,716	57.8%	82.5%

Bottom line: after including the 7.2M tokens Loom itself used, total token consumption was an estimated 57.8%-82.5% lower than the same work would have required without Loom. Each token Loom used corresponded to an estimated 16.6-54.7 tokens of avoided work.

Where the token savings came from

These categories divide the same total savings shown on page 1. They are not additional savings.

Type of saving	How Loom saves tokens	Conservative	Realistic
Reusing known facts and instructions	Provides information and instructions that were already established, avoiding repeated code searches, history searches, experiments, or questions to the user.	5,939,295	17,644,658
Preventing unsafe or unapproved actions	Reminds the model of important decisions and limits before it acts, preventing unsafe, unapproved, postponed, or previously rejected work.	2,136,570	10,140,322
Avoiding duplicate, outdated, or wrong work	Shows what was already completed, disproven, or replaced, avoiding repeated investigation and implementation based on old assumptions.	17,917,177	66,593,375
Preventing release and multi-computer mistakes	Keeps release requirements and computer-specific test evidence available, preventing incomplete releases, false approvals, repeated diagnosis, and platform failures.	55,771,770	177,455,635
Avoiding invalid measurements and reruns	Keeps study definitions, procedures, known defects, and comparison rules available, avoiding invalid analysis and expensive reruns.	38,026,206	121,683,864
Total		119,791,018	393,517,854

How the estimate was produced

1	Start from actual use. The study examined real Codex interactive sessions, Codex automated sessions, and Claude automated sessions.
2	Replay representative cases without Loom's saved context. Controlled comparison runs measured whether the model searched code and history, pursued unnecessary work, or stopped to ask the user.
3	Apply the measured results only to similar cases. Related Loom lookups that supported one decision were counted once, preventing the same benefit from being counted multiple times.

How to read the estimates. The conservative estimate values avoided work cautiously and limits the credit assigned to each decision. The realistic estimate also includes the hours or days that serious release, implementation, and measurement mistakes can consume. The figures measure tokens supplied to the models during the study period; they are not a company-wide projection.

Two comparisons are reported. Total-token reduction = (estimated tokens without Loom - tokens used with Loom) / estimated tokens without Loom. The estimate includes Loom's own token use. Savings per Loom token = estimated avoided-work tokens / tokens used by Loom. Both use the same observed sessions and the same estimated savings.