

AI Use Policy & Disclosure



Responsible AI. Real Impact.

We use AI thoughtfully and transparently to support our mission, protect our values, and create lasting environmental impact.



Responsible by Design

We prioritize ethics, accuracy, and fairness in our AI tools.



Data Privacy & Security

We protect sensitive information and follow industry best practices.



Human Oversight

AI supports our work—people guide our decisions.



Purpose-Driven Innovation

We use AI to advance forest health, carbon solutions, and a more sustainable future.



ROOTED IN NATURE.
COMMITTED TO RESPONSIBLE INNOVATION.

Transparency today. A better tomorrow.

This document explains how we use AI in the work we deliver to partners and clients across Sherman Ventures & Advisory, Allied Earth Biocarbon, Tsimimalanoxw LP, and CAPSA Carbon. Wherever it makes sense, we look to leverage AI responsibly and transparently - we think that is the right thing to do in today's evolving world.

Not long ago, getting work like this done well meant assembling a large team of experts and aligning them on a problem. Today, we have access to remarkably powerful tools for research, brainstorming, background checking, and the kinds of tasks that were previously slow and costly to do well. The art of the work has shifted - but the standards have not.

British Columbia's Professional Governance Act connects most regulated professionals across the province to a common framework of professional standards, conduct, and accountability. Our principal, Kelly Sherman, is a professional forester working under that framework and serves as our companies' certified corporate lead under EGBC (Engineers and Geoscientists BC), having completed the corporate lead training. We are clear-eyed about what it means to hold professional credentials, where our expertise lies, and where it does not. We take the codes of ethics that govern professionals seriously, and we value real expertise.

The principle running through what follows is simple: AI helps us think and draft; we decide what is right. Everything below describes what we actually do, not an aspiration.

The tools we use

We were early adopters of ChatGPT - within days of its initial release - and have kept up to speed and tested many AI tools since. We have settled on a small, deliberate stack that serves us well at this point in time. Like any technology this will evolve, but in May 2026 the working setup is as follows.

Tool	Role	Where it runs
Claude	Core working assistant	Cloud (Anthropic)
ChatGPT	Strategy reviewer / cross-check	Cloud (OpenAI)
LM Studio (local model)	Privacy-protected workspace	Local on our laptop — no internet transmission

Claude is our primary AI. Each engagement is set up as a project in Claude, and Claude is connected to Asana, where we manage our to-do list and priorities. We find Claude particularly strong for document management and production.

ChatGPT is the tool we have the most experience with, and we still use it primarily for strategic review. We set up a parallel, similar project in ChatGPT and use it to review Claude's recommendations or help with more complex matters. Specific strategies sometimes get bounced between the two dozens of times, until both reach a strong, consistent answer that we are aligned on.

LM Studio runs a language model locally on our laptop. We use it for any work that cannot be safely de-identified — its job is privacy, not strategy.

Confidentiality and how we decide what goes where

Our default is to de-identify material before any cloud-based tool sees it. If a piece of work cannot be safely de-identified - because the identity of the client, organisation, or individual is intrinsic to what we are doing - it stays on the local model. The cloud tools never see your name attached to your situation.

In practice that means, before anything goes into Claude or ChatGPT:

- Names of clients, organisations, partners, and individuals are removed or replaced with generic placeholders.
- Specific identifying details (locations, products, dates that pin down who is involved) are stripped or generalised.
- Commercially sensitive figures and unreleased information are removed unless they are essential and can be abstracted.
- Anything we would not be comfortable having read aloud in a public setting is treated as needing local handling.

When the material is genuinely confidential and cannot be de-identified without losing what makes it the question, we use LM Studio, which runs entirely on our laptop. Inputs and outputs are not transmitted to any third party.

We do not enable training on our data in any cloud AI tool we use, and we review those settings periodically.

What stays human

Some things we do not delegate to AI, and never will:

- Judgement on what to recommend. AI helps us think through options; we decide what to advise.
- The final read of anything sent to you. Nothing leaves our desk without us having owned the words and the conclusions.
- Quotes, citations, statistics, and named claims that appear in deliverables. Where AI suggests these, we treat them as leads to verify, not facts to repeat.
- Communications in our voice. AI does not write things and send them as if from us.

Whichever tools we use to produce a piece of work, the work is ours. We are accountable for accuracy, judgement, and outcomes — not the tools.

How we check the work

We want to be honest about what cross-model checking does and does not do, because the difference matters.

What it does well:

- Surfaces reasoning gaps, internal contradictions, and counter-arguments we had not considered.
- Catches conclusions that do not follow from their premises.
- Provides a useful second view on directional or strategic questions — if Claude and ChatGPT disagree, that is a signal an area needs more care.

What it does not reliably do:

- Verify that a citation, quote, or specific statistic is real. Two AI tools can confidently agree on something that neither has actually seen, so cross-model agreement is not source verification.
- Replace specialist domain knowledge.

Our substantive check is our own expertise in the area we are working in. We notice when AI output is wrong because we know the subject. Where work depends on a specific source, quote, or figure, we confirm it against the original — cross-model checking is not how that gets done.

How we disclose AI use

Our default is to disclose, not to leave it for you to ask. Deliverables we produce will carry a short standing line letting you know AI was part of the process and clarifying what was not delegated to it.

Sample standing disclosure line

"Developed with the assistance of Claude and ChatGPT. Confidential material was handled on a local model. Final analysis and recommendations are our own."

If you want a more detailed account of how AI featured in a specific piece of work — which tool was used for what, what stayed local, what we verified by hand — we are happy to walk you through it. We would rather have the conversation than leave it to assumption.

A note to close

We've written this to be useful in real engagements — a practical reference for how AI shows up in our work, and a starting point for any conversation you'd like to have about it. If anything in it is unclear, or if you have a specific concern about how AI should or should not be used on an engagement we are working on, please raise it directly. We will adapt our approach to the context of our work together, and we will tell you when we have.

—

Kelly Sherman

Principal — on behalf of Sherman Ventures & Advisory, Allied Earth Biocarbon, Tsimimalanoxw LP, and CAPSA Carbon

kelly.l.sherman19@hotmail.com