

eBook

AI

ESSENTIALS FOR DLGG

**FIRST
EDITION**

MATTHEW DOBSON

AI Essentials for Dominion Lending Centres Inc. All rights reserved.

No part of this eBook may be reproduced, shared, or distributed in any manner whatsoever outside the terms of the applicable license agreement, without the express written permission of the publisher, except for the use of brief quotations in a book review.

This eBook is licensed for use only by Dominion Lending Centres Inc. and its authorized personnel, and may not be distributed, copied, or transmitted to any other person or entity without prior written consent. For additional permissions, please contact the publisher.

AVA Credit Inc., 582 St Clair Ave W Unit 302, Toronto, Ontario, M6C 1A6, <https://avacredit.ca>

Copyright © 2025 by AVA Credit Inc.

First Edition

Table of Contents

| | |
|--|----|
| About this eBook | 4 |
| Introduction | 5 |
| Chapter 1: Building Blocks | 7 |
| Chapter 2: Inside the Engine | 9 |
| Chapter 3: The RGI Framework for Human-AI Collaboration | 10 |
| Chapter 4: Prompting Powers the Framework | 12 |
| Chapter 5: Collaborating with Generative AI in Practice | 14 |
| Chapter 6: Benefits and Risks of Using Generative AI | 15 |
| Chapter 7: Regulations and Governance | 19 |
| Chapter 8: The Do's & Don'ts, Responsible Use and Guardrails | 20 |
| Appendix: Quick AI Essentials Guide | 22 |
| Appendix: Important Disclaimer | 23 |

About this eBook

This guide fills a critical gap in the AI landscape by offering a universal, practical framework for working with generative AI across different roles, tools, and workflows. Instead of simply cataloging use cases, it introduces the RGI Framework™, which stands for Retrieve, Generate, and Interact. This framework distills the complexity of large language models into a clear, repeatable methodology that enables effective collaboration between professionals and AI.

Developed for mortgage professionals at Dominion Lending Centres Inc, Canada's largest mortgage brokerage network, the guide moves beyond experimentation into precision. It helps users answer not just "What can AI do?" but "How can I use AI to retrieve, generate, or interact with my information and data for real results?" Grounded in first principles, it empowers professionals to become skilled operators who work with intent and clarity. This foundational playbook delivers high-impact, responsible adoption of AI language models in financial services.

How to Use this eBook

This eBook serves two purposes. You can read it traditionally to build expertise in human-AI collaboration or upload it to an enterprise-grade AI model to create an interactive advisor that can answer questions, clarify concepts, and guide implementation using the RGI Framework™ through semantic search. Both approaches unlock the practical knowledge needed to work effectively with generative AI in financial services. Whether you choose linear reading or conversational querying, you'll gain the insights necessary to succeed.

About the Author

Matthew Dobson is a seasoned financial services leader in Toronto with more than two decades of expertise in commercial lending. He is the founder of AVA Credit Inc., a pioneering platform that integrates generative AI into mortgage and business finance workflows. He has held senior sales and credit leadership roles at GE Capital and TD Bank, and he served as Chief Credit Officer for a bank-backed private credit fund in Los Angeles, specializing in complex commercial financing. Over his career, he has originated and underwritten billions of dollars in commercial mortgages and equipment loans across Canada and the United States.

In 2019, Matthew trademarked AI Credit® in Canada and pioneered the use of generative AI models in live credit underwriting. From 2020 to 2023, he led AI-enabled credit trials on U.S. Small Business Administration bank-grade mortgage transactions, demonstrating how human expertise and large language models can collaborate to improve lending decisions. A landmark transaction in 2021 involved a multimillion-dollar loan to a data center supporting early AI training infrastructure for a leading foundation model, showcasing his deep understanding of AI and his ability to bridge emerging technologies with institutional finance.

Selected as a 2025 Perplexity AI Business Fellow, Matthew combines deep financial acumen with hands-on operational expertise with AI to drive innovation in lending and enterprise workflows. Born and raised in Toronto, he is a graduate of Western University.

Introduction

Why AI Matters for Mortgage Professionals

The mortgage industry runs on relationships, judgment, and trust. These fundamentals haven't changed. What has changed is the speed and scale at which professionals need to operate. Client expectations have accelerated. Competition has intensified.

Consider what a typical mortgage agent handles daily: researching market conditions, drafting client emails, creating presentation materials, summarizing complex borrower situations, and managing follow-up communications. Each task requires expertise, but much of the execution is routine. This is where generative AI transforms the game.

Generative AI lets humans use natural language to communicate with advanced computing systems to retrieve, generate, and interact with the knowledge, content, and decisions needed to get work done.

Generative AI doesn't replace the human elements that drive mortgage success. Instead, it amplifies them. Like the calculator for accountants or word processors for writers, AI tools handle the mechanical work so professionals can focus on what matters most: building relationships, making sound judgments, and serving clients.

Need to research a local real estate brokerage before a meeting? AI can scan public data and generate a conversation agenda. Writing nurture emails for first-time buyers? AI becomes your copywriter, maintaining your voice while saving hours. Explaining a complex borrower story to a lender? AI helps distill the situation into clear, persuasive language.

The value isn't in what AI does alone. The real value lies in what it enables you to do better. More time for client relationships. Faster response times. More consistent communication. Deeper market analysis. These improvements compound over time, creating sustainable competitive advantages.

The RGI Framework: Your Guide to AI Partnership

Most professionals struggle with AI because they lack a mental model for how it works. Vendor demonstrations show flashy features but often miss the core fundamentals. Technical articles dive too deep into mechanics that don't matter for daily use.

This guide introduces the RGI Framework, a first principles approach to understanding and applying AI in mortgage workflows:

- **Retrieve:** Ask AI to find, search, and extract information from knowledge sources
- **Generate:** Use AI to write, create, and develop content from emails to presentations
- **Interact:** Engage with AI as a thinking partner and workflow assistant

Every use case for AI fits within this framework. Master these three capabilities, and you can navigate any generative AI tool or workflow without getting lost in technical complexity.

Partnership, Not Replacement

AI is a co-pilot, not an autopilot. You remain the licensed professional making decisions, maintaining client relationships, and ensuring compliance. AI provides speed, consistency, and processing power. You provide judgment, context, and expertise.

This partnership approach matters especially in regulated industries like mortgage lending, where accuracy and accountability are non-negotiable. AI outputs require human validation. Client interactions demand personal touch. Compliance decisions need professional oversight. The most successful mortgage professionals will be those who learn to collaborate effectively with AI tools while maintaining the human elements that clients value most.

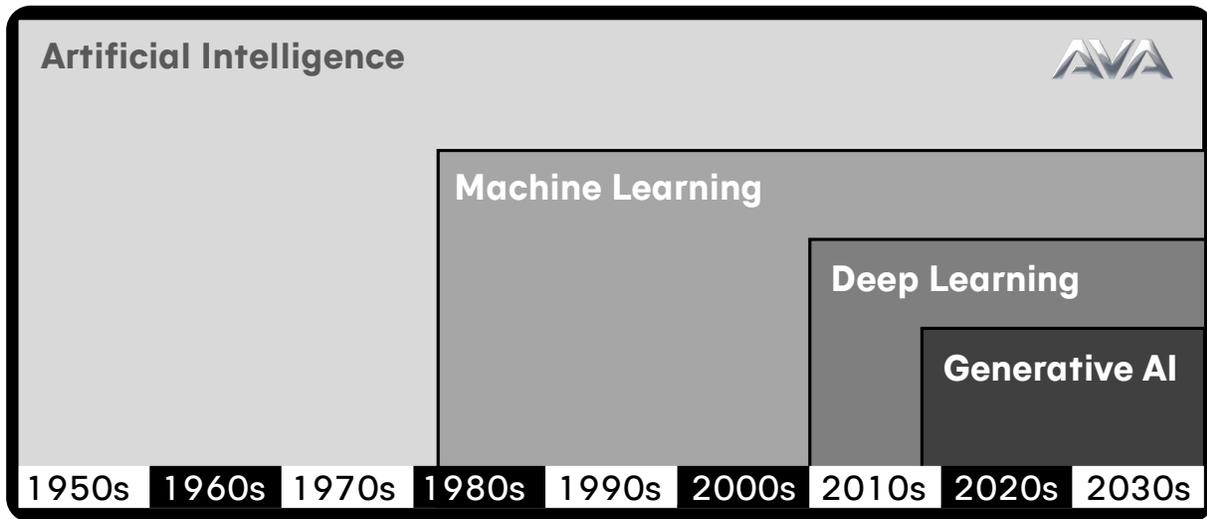
What You'll Learn

This guide will help you:

- ✓ **Understand the fundamentals**
Explain large language models in plain English so you understand what they are, how they work, and why they matter today.
- ✓ **Apply the framework**
Align the RGI framework with your workflow to surface and launch high-impact generative AI use cases in your daily tasks.
- ✓ **Master practical skills**
Cultivate practical prompting habits that power each outcome, activate the RGI framework, and enable the human-AI partnership.
- ✓ **Maintain professional standards**
Set out guardrails and governance because accuracy, privacy, and compliance are non-negotiable in financial services.

Let's begin.

Chapter 1: Building Blocks



Artificial Intelligence

Any system that pursues a goal by processing data in a way that mimics human judgment, without step-by-step human rules. AI swaps hard-coded “if-this-then-that” logic for software that adapts.

Algorithm

Step-by-step instructions or rules a computer follows to solve problems or perform tasks.

Machine Learning

Branch of AI focused on building systems that learn from data rather than relying on explicit programming, making software more accurate and responsive over time.

Deep Learning

Powerful subset of machine learning that uses layered artificial neural networks to identify complex patterns in large, high-dimensional datasets. As neural networks grow in depth and scale, so does their capacity to interpret complex information.

Generative AI

AI systems designed to create new, original content such as text, images, audio, or code by learning patterns from existing data and then generating context-appropriate outputs on demand, rather than merely analyzing or labeling data.

Foundation Models

Large-scale neural networks that power today’s generative AI landscape. Their emergence began with the 2017 introduction of the transformer architecture, detailed in the landmark paper “Attention Is All You Need,” which redefined how machines process language. The landscape has evolved rapidly from text-only models to sophisticated multimodal systems capable of processing and generating content across text, images, audio, video, code, and even specialized domains like protein structures and robotic control. Their scale, flexibility, and transferability make them the critical layer beneath the generative AI ecosystem, allowing professionals to collaborate with intelligence at speed and scale.

Large Language Models

Foundation Models trained chiefly on text and optimised to read and write human language. Using the transformer’s self-attention mechanism, they predict one token at a time across billions of parameters, yielding fluid prose, sharp summaries, multilingual translation, and conversational Q&A from a single prompt. In practice, an LLM is the engine behind most chat interfaces, writing assistants and enterprise systems in use today.

What Made Generative AI Possible?

Generative AI refers to artificial intelligence systems that can create new content, such as text, images, or code, rather than simply analyzing existing data. Its rise has been driven by three key breakthroughs:

First, the invention of the Transformer architecture in 2017, which radically improved how machines understand and generate language by enabling models to capture context more effectively.

Second, an explosion of digital content, from websites and e-books to source code and corporate filings, created the vast training data needed to teach these systems how to write, reason, and respond.

Third, advances in computing power, especially the rise of specialized chips like GPUs and TPUs, made it feasible to train and deploy large-scale models at speed and scale.

Together, these three pillars have opened the door to powerful AI tools that financial professionals can now apply across research, communication, analysis, and client service.

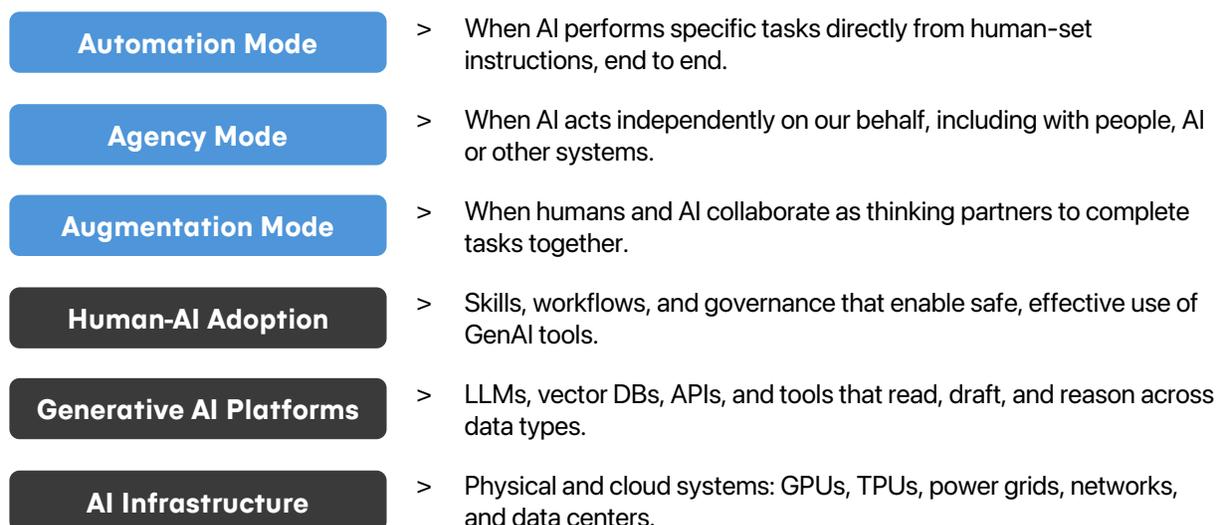
The Generative-AI Tech Stack

Generative AI isn't a single tool. It's a layered stack. Each layer builds on the one below, creating the infrastructure behind modern AI applications. At the base are hardware and cloud systems: GPUs, TPUs, data centers, and high-speed networks. Above that are GenAI platforms, including LLMs, vector databases, orchestration tools, and secure APIs that read, reason, and generate across data types. Sitting atop the stack is the human-AI adoption layer, where skills, workflows, governance, and culture shape real-world impact. This layer enables three modes of interaction: Augmentation (today's co-pilot), Agency (emerging assistants that act on our behalf), and Automation (fully autonomous execution).

The Financial Services Stack

Financial services work is built on judgment, trust, and regulation, three areas where a human in the loop is non-negotiable. In regulated, relationship-driven fields like banking, lending, capital markets, and accounting, the near-term opportunity lies in the Augmentation layer: an AI co-pilot handles the heavy lifting while you apply judgment, ensure compliance, and manage client nuance. Real enterprise value comes not from models alone, but from how humans adopt and apply them.

Generative AI Architecture: Core Stack Layers and Human-AI Collaboration Modes



Chapter 2: Inside the Engine

Mortgage professionals are not expected to code or train large language models, but a basic understanding of how they work is now part of modern financial literacy. Knowing what happens under the hood helps you judge when an LLM is the right tool, craft precise prompts that deliver relevant output, and explain results clearly to clients, colleagues, or compliance teams. In short, this is mechanics for operators, not a deep dive for engineers.

An algorithm is a set of step-by-step instructions a computer follows to solve a problem or perform a task. In generative AI, algorithms train large language models like ChatGPT by guiding them to detect patterns in data, adjust internal settings, and learn how to generate useful responses. If the model is the engine, the algorithm is the blueprint and tuning process used to build and optimize that engine. Mortgage professionals don't need to know how to write algorithms, but understanding that they exist helps demystify how these systems work and why they're so powerful.

Every Large Language Model follows the same three-stage cycle:



During training, the model reads billions of tokens drawn from books, filings, websites, code, and more. By repeatedly guessing the next token and comparing its guess to the correct answer, it learns the statistical rhythms of language—a method called self-supervised learning. A second pass, often guided by human feedback, fine-tunes the model so its answers line up better with real-world intentions. Training is where the model acquires its broad knowledge.

When you type a prompt, the model moves to tokenization. Your text is chopped into sub-word units called tokens, and the model processes them inside a fixed context window. That window defines how much information the system can “see” at one time and therefore shapes cost, speed, and output length. A 50-page loan agreement may need to be chunked or augmented with retrieval tools because its token count can exceed the window.

Finally, the model enters inference. It calculates the most likely next token, appends it, recalculates, and continues—literally word by word—until it decides to stop. Settings such as temperature and top-k sampling govern how creative or conservative this process feels: low temperature for drafting policy language, higher for brainstorming marketing copy. Different decoding strategies explain why the same model can produce a tight covenant summary in one moment and a free-flowing thought piece the next.

Understanding these three stages turns a large language model from a black box into a usable, predictable tool. Each part of the RGI Framework connects directly: **Retrieve** tasks are limited by the context window, **Generate** tasks respond to decoding settings, and **Interact** tasks depend on how clearly you prompt and refine.

Chapter 3: The RGI Framework for Human-AI Collaboration

Generative AI enables humans to use natural language to communicate with advanced computing systems, unlocking three core capabilities that form the foundation of human-AI collaboration: Retrieve, Generate, and Interact. This RGI Framework serves as the first principles of AI adoption, capturing how professionals derive value from language models by retrieving knowledge, generating content and ideas at scale, and interacting with co-pilots, agents, systems, documents, and workflows through natural conversation. Rather than learning complex technical interfaces, professionals can now accomplish sophisticated work simply by communicating their needs in plain language.

RGI Framework™

RGI is the first principles framework for human-AI collaboration.

Retrieve: Ask • Find • Search
Generate: Write • Ideate • Create
Interact: Instruct • Query • Collaborate

Whether you are a financial professional working with AI on the front lines of your trade, or an enterprise leader designing generative AI systems for scale, every meaningful use case can be understood through this lens. This is first principles thinking.



When we **retrieve**, we use language to extract specific information from documents, databases, and historical records instantly and without manual effort. Generative AI's first super-power is friction-free search. The model ingests your prompt, then pulls facts from its training, a live retrieval layer, or both. LLMs eliminate the need to search line by line, surfacing accurate and contextual responses from both internal and external knowledge sources.

✓ *Example, "I've uploaded a summary document comparing high ratio vs. conventional mortgage insurance rules. Pull the key differences in down payment requirements"*

When we **generate**, we accelerate the mechanical work of writing, planning, brainstorming, and content creation. Whether drafting marketing copy, credit memos, investor decks, client emails, or generating code, imagery, or solutions from scratch, LLMs enable faster output, richer variation, and stronger first drafts. This allows professionals to explore options, develop content, test strategies, fail fast and reduce time spent on repetitive tasks.

✓ *Example, "Write 10 social media posts that guide a first-time homebuyer in the GTA through each step of the mortgage journey—from pre-approved to closing."*

When we **interact**, we engage with a smart, task-oriented thinking partner, workflow assistant, or orchestrator. Financial professionals can chat with documents, extract structured data, pre-fill forms, and guide multi-step workflows through natural conversation. Prompting becomes a live dialog of follow-ups, clarifications, and reasoning.

✓ *Example, "Walk me through the steps to get a client from pre-approval to funding, including tips, watchouts and follow-up questions for keeping the deal on track."*

RGI is more than a way to describe how large language models deliver value. It is a practical framework for understanding what they do, why it matters, and how to apply them effectively. Whether you're making real-time decisions on the front lines or designing enterprise systems at scale, RGI clarifies both the purpose behind the tool and the mechanics behind the results.

This isn't about replacing people. It's about combining human expertise with AI precision to sharpen judgment, accelerate execution, and free up time for what matters most: strategy, creativity, and meaningful relationships.

Whether you're working with a general-purpose public model, a custom enterprise deployment, or a co-pilot powered by retrieval-augmented generation, the RGI framework holds. The tools may change. But the principles, and the prompting practices that activate them, remain the same.

Thousands of new tools and architectures will emerge. But at the center of it all will still be a human and a language model, working together to retrieve, generate, and interact. RGI is the foundation for that collaboration. It applies universally: from frontline professionals to C-suite leaders designing secure, AI-enabled workflows.

Strip away the vendor hype and technical jargon, and what remains is this: professionals asking better questions, generating better output, and interacting more efficiently with intelligent systems. That's the shift.

Why First Principles Win

Most enterprise software automates isolated tasks. RGI reflects how professionals actually work: by retrieving information, generating content, and interacting with systems. Apply this framework, and the 75-year arc from theoretical AI to practical tools collapses into a repeatable workflow:

- ✓ You ask.
- ✓ The model retrieves, generates, and interacts.
- ✓ You apply professional judgment and move forward.

That's it. A new workflow for the AI era. Grounded in clarity, led by human judgment, and ready for real work.

The key to unlocking each capability is prompting. In the next chapter, we'll break down how to structure prompts that activate the RGI framework with precision.

Chapter 4: Prompting Powers the Framework

If RGI is the foundational framework that explains what large language models can do, prompting is the core skill that makes it work. Prompting is how professionals communicate with generative AI by using clear, structured language to express intent, define outcomes, and shape the response. It's a core capability that determines the quality of every interaction with AI.



Effective prompting begins with a **GOAL** stated clearly. The AI cannot guess what you mean, so you must tell it explicitly. Whether you're asking for a summary, a draft, or a recommendation, be direct and specific. Next, always ensure that **CONTEXT** is given. Frame the situation by explaining who the audience is, what they need to know, and why it matters. AI models are highly responsive to context, and the more relevant detail you provide, the better the result.

Once your goal and context are in place, **FORMAT** should be defined. If you need bullet points, a table, or a paragraph with a certain tone, say so. Don't leave formatting to chance. Adding **REFERENCE** materials sharpens the output further. Whether it's an uploaded PDF, a data sheet, or text from a previous message, citing reliable content helps the AI stay grounded and accurate.

Prompting is not a one-and-done task. It's a back-and-forth process. If the first draft isn't quite right, the prompt should be **REFINED** and iterated. Ask for a shorter version, a different tone, or a new format. Iterate intentionally. This is how professionals get high-quality responses and avoid generic, vague, or hallucinated output.

To help professionals internalize great prompting habits, this guide outlines a simple pattern called the Prompt Formula. Think of it as a modular sentence template you can adapt to any situation by plugging in the right variables. This formula combines all five best practices into a single, actionable flow.

Prompt Formula

Start by clearly stating your goal. Provide relevant context, including who the audience is and why the task matters. Specify the desired output format and tone. Cite any references or source material to ground the response. Include any constraints or special instructions as needed. After reviewing the response, refine and iterate your prompt to improve results.

Prompt Checklist

| Element | Complete? |
|-------------------------|-----------|
| Clear Goal Stated | ✓ |
| Relevant Context Given | ✓ |
| Output Format Defined | ✓ |
| Reference Docs or Links | ✓ |
| Prompt Refined | ✓ |

Prompting AI is a strategic skill. The best results come from structured input, guided iteration, and thoughtful use of the model's strengths.

Power Prompting Like A Pro: Techniques for Advanced Users

Once you've mastered the basics, you can apply more sophisticated techniques to control quality, tone, and depth of response. These are used by top operators across finance, law, and consulting.

✓ **Think-First Prompting (a.k.a. Chain-of-Thought Reasoning)**

Ask the model to walk through its logic before giving a final answer. This reduces errors, reveals assumptions, and creates stronger reasoning — especially useful for deal memos, underwriting decisions, or policy questions.

Example, "I'm preparing a webinar for first-time homebuyers. Walk through the key topics they care about most, then suggest a clear order to present them in."

✓ **Role and Persona Prompting**

Assign a specific role to steer the tone, depth, and vocabulary. This helps ensure output is tailored to your audience — a client, committee, regulator, or investor.

Example, "Respond as a licensed mortgage agent explaining a refinance strategy to a first-time homebuyer. Keep it calm, clear, and reassuring."

✓ **Break Complex Tasks into Sequential Prompts**

Guide the AI's reasoning by breaking down the task into clear, ordered steps and delivering them as sequential prompts. This improves accuracy, logic, and the actionability of the response.

Example: "Help me create a lead generation campaign targeting Vancouver condo investors. First, list key pain points. Next, draft a cold outreach email. Then suggest three follow-up questions to ask during a phone call."

✓ **Secret Weapon: Ask the Model for Prompting Advice**

You can ask the model how to prompt it — and it will tell you.

Example, "I'm trying to write better client update emails. What's the best way to prompt you to help with that?"

Common Pitfalls to Avoid

- ✗ Vague prompts – No specific outcome or unclear intent
- ✗ Overloaded asks – Multiple goals in one prompt
- ✗ No formatting guidance – Leaves structure to chance
- ✗ Lack of iteration – Accepting weak or generic output
- ✗ Missing context – Forces the AI to guess (accuracy drops)

Chapter 5: Collaborating with Generative AI in Practice

Generative AI is not an automation tool. It is an amplifier of inherent human abilities and a collaborator that relies on human judgment, intent, and contextual understanding to be effective. The large language model contributes speed, scale, precision, and computational prowess. When combined, this human-AI partnership produces work that is accelerated, sharper, and more thoughtful. Success does not come from issuing casual commands but from engaging in a structured and deliberate exchange. The model cannot infer intent unless that intent is expressed with precision. This is why effective prompting is not a technical trick or a passing trend. It is a core skill for the foreseeable future.

Large language models are exceptional at processing large volumes of information, generating content variations quickly, spotting patterns in prior data, and adapting instantly to structured instructions. But they lack judgment, context, and discretion. These are responsibilities of the human operator. It is the operator who provides direction by shaping the prompt, embedding relevant business or legal context, and exercising discretion in accepting, refining, or discarding the model's output. The human operator is also the final safeguard for ensuring accuracy, ethical integrity, and domain-specific appropriateness. The quality of the result depends as much on the clarity of the human input as it does on the model's technical capacity.

Working with a large language model is not automation. It is collaboration with a high-speed reasoning engine that responds to the exact language provided. The difference between a vague, unusable answer and a sharp, production-ready result often comes down to how clearly the prompt was structured. The tool remains constant. The outcome is shaped by the skill of the operator.

Just as a trained pilot, surgeon, or crane operator consistently outperforms an untrained observer, professionals who know how to collaborate with LLMs gain a significant advantage. Prompting a language model isn't about shortcuts or clever tricks; it is about communicating with precision, providing structure and context, and refining inputs to achieve superior results. Prompting serves as the essential bridge between human intent and machine capability.

As an enduring message to mortgage professionals and enterprise leaders, the value of the human operator is not diminished by AI. It is elevated and amplified, but only when the human remains actively engaged with the system. The future of work is not a choice between human or machine; it will be defined by the partnership between the two. In this evolving environment, success depends on treating AI not as a black box or a magic tool, but as a collaborator that requires human judgment, input, and oversight. Generative AI is a co-pilot. It is not an autopilot.

Chapter 6: Benefits and Risks of Using Generative AI

Before diving into the details, here’s a quick snapshot of what Generative AI brings to mortgage professionals: the powerful benefits on one side, and the equally important risks and burdens on the other. Generative AI is a force multiplier, but it comes with trade-offs that must be managed carefully.

| Benefits of Using Generative AI | Risks of Using Generative AI |
|--|---|
| <p>Accelerates routine tasks Turns hours of work into minutes by retrieving, drafting, and summarizing instantly.</p> <p>Frees time for higher-value work Shifts focus from repetitive tasks to strategy, creativity, and client relationships.</p> <p>Scales expertise Acts as a tireless junior analyst, writer, and reviewer available 24/7.</p> <p>Accelerates learning Rapidly ramps up new or junior staff by delivering instant, context-specific answers.</p> <p>Interact with internal knowledge Access and engage with internal knowledge, including policies, training materials, and best practices.</p> <p>Enhances decision-making Helps refine strategies, stress-test scenarios, and surface hidden insights.</p> <p>Codifies and scales workflows Reusable prompts multiply efficiency across teams and departments.</p> <p>Fuels collective learning Sharing prompts and best practices builds stronger teams and communities.</p> | <p>Compliance & security exposure Public models can leak sensitive data and create legal liabilities.</p> <p>Hallucinations Confident but false or misleading outputs require careful fact-checking and human validation. Improved prompting habits also help.</p> <p>Decision fatigue & creative flattening Overuse can lead to repetitive thinking and reduced originality.</p> <p>Over-reliance on AI Can erode human judgment and weaken meaningful interpersonal interactions.</p> <p>Bias & fairness risks Models reflect data biases, requiring careful review for ethical and regulatory compliance.</p> <p>Cognitive overload Sustained prompting, reviewing, and refining hundreds of outputs can be mentally draining.</p> <p>Shadow AI Unsanctioned use of AI tools can bypass IT and compliance safeguards, creating hidden risks.</p> <p>Systemic risks Organizational over-dependence on AI can degrade trust, governance, and long-term resilience.</p> |

The Benefits of Human-AI Collaboration

Working with a large language model is not just an upgrade in tooling. It is a shift in tempo, scale, and cognitive engagement. For professionals who bring deep expertise to high-stakes environments such as financial services, law, accounting, consulting, healthcare or enterprise operations, the benefits are profound. Language models compress the distance between intent and execution, turning hours of work into minutes. They retrieve critical information from documents in seconds, generate high-quality drafts on demand, and interact seamlessly across email, spreadsheets, CRMs, and internal workflows. This ability to retrieve, generate, and interact with speed and precision enables professionals to ideate faster, iterate more often, and make better-informed decisions.

The most profound benefit of human-AI partnership is not just accelerating routine tasks or amplifying expertise. It is creating time for higher-value activities such as building meaningful client relationships, strengthening trust with colleagues, and focusing on the creative, strategic, and interpersonal work that drives real value. By dramatically increasing the speed and efficiency of information retrieval, content generation, and workflow execution, LLMs free professionals from routine, repetitive tasks. This reclaimed time is not incidental; it is the essential dividend of AI augmentation.

The capabilities described in the RGI framework, which include retrieving knowledge, generating content, and interacting with systems, are not theoretical concepts; they translate into measurable gains in real-world execution. LLMs can serve as tireless junior analysts, writing partners, or reviewers, available around the clock and unaffected by fatigue. Whether reviewing contracts, analyzing credit memos, drafting client communications, or learning to build or use a new technology application, these models help operators work with greater agility and control. Prompting acts as a skill multiplier, transforming a professional's intent into structured output with customizable tone, format, and structure. Over time, professionals codify their workflows, developing reusable prompts and processes that scale their expertise across departments or teams. The result is amplified output, accelerated learning cycles, and the ability to test ideas, fail fast, and refine strategies in real time.

One of the most transformational benefits of enterprise-grade large language models is their ability to retrieve and interact with internal company knowledge, including: training manuals, product documentation, marketing assets, HR policies, compliance protocols, and other proprietary content stored across intranets and cloud drives. When tuned with this organizational context, a language model becomes more than a tool; it becomes a 24/7 AI co-pilot and mentor. For new or junior employees, this dramatically shortens the ramp-up time from novice to expert by delivering instant, reliable answers rooted in company-specific practices.

For experienced professionals, the impact is just as significant. Language models enhance expertise by enabling rapid access to institutional knowledge and simulating high-stakes scenarios, such as difficult client conversations or complex compliance reviews. These systems help professionals refine strategies, stress-test responses, and stay current on internal best practices. The result is a workforce that learns faster, performs smarter, and scales knowledge organization-wide with unprecedented speed and precision.

The Hidden Burdens on Operators

But these advantages do not come without cost. The burdens of working with advanced AI systems are subtle, cumulative, and often underestimated. Cognitive overload is the most immediate. While the model generates output instantly, the responsibility to craft effective prompts, review each result, spot hallucinations, and iterate toward the right answer lies squarely with the human operator. This level of sustained mental engagement, repeated over hundreds of interactions, can become draining. When combined with the pressure of high-stakes decision-making, the effect is amplified.

As the AI infrastructure and generative AI platforms improve in the months and years ahead, and as the more powerful the system becomes, the more critical the human oversight. Without thoughtful design, prompts and responses may become repetitive, leading to creative flattening and decision fatigue. Perhaps the most overlooked pressure point in the system is the human operator.

Operational and Systemic Risks

Hallucinations, meaning confident but false or misleading outputs, remain a persistent risk. These issues are not always caused by a flawed model; more often, they result from unclear prompts or missing context. Language models do not evaluate truth or understand consequences. They generate text by predicting what is most likely to come next, based on patterns in their training data. Professionals must treat every output as a draft, not a decision, and apply critical thinking before acting on AI-generated content. In domains like financial services, credit underwriting, legal review, or compliance, the risk of relying on unverified outputs is not just operational. It can have serious consequences.

Other risks of using LLMs are more systemic, meaning they arise from the broader patterns and structures within organizations or workflows rather than from individual actions. Over-reliance on AI can gradually erode human judgment, especially when professionals begin to defer too often to machine-generated output. This dependence can also distance us from meaningful human interaction, weaken relationships, and diminish the value of community-based learning, all of which are essential for professional growth and well-being.

Bias, Compliance, and Information Security

Further, there are the non-negotiable risks of bias, compliance, and information security. Models trained on large datasets inevitably reflect the biases present in those datasets. Outputs must always be reviewed for fairness, accuracy, and ethical soundness. In regulated industries, AI-generated content must meet legal and policy thresholds, and enterprise-grade controls must be in place to protect sensitive data. Public models and unsecured APIs may offer convenience, but they also carry significant exposure risks for sensitive, personal, or confidential information.

A rapidly emerging risk in financial services is the rise of “shadow AI.” In a bank setting, shadow AI refers to employees using non-bank-approved AI tools for work tasks, bypassing enterprise IT and compliance safeguards. In a large mortgage brokerage, the risk takes a different but equally serious form. Here, shadow AI occurs when mortgage agents, who are independent contractors rather than employees, use free or low-grade public models for lender-facing workflows, fail to mask client-sensitive data, or co-mingle personal and business accounts. These practices bypass the brokerage’s enterprise-grade protections, creating risks of data leakage, privacy breaches, and potential governance gaps that can jeopardize both lender relationships and an agent’s professional standing.

Shadow AI undermines data governance, exposes confidential information, and can lead to breaches of privacy laws. Proactive governance, clear policies, and deliberate agent education are essential to prevent these risks and ensure safe, auditable AI adoption across the network.

Community, Vigilance, and Shared Responsibility

The benefits and risks of large language models are shaped by the human communities that use them. Generative AI delivers the most value when professionals share insights, prompt strategies, and best practices, fostering collective learning and innovation. At the same time, systemic risks are best managed through shared vigilance and collaborative problem-solving. Building and participating in a community of users is the most effective way to maximize AI’s advantages while addressing its challenges.

Working with a language model is a game-changer, enabling professionals to focus on complex problem-solving and meaningful client relationships. But this power also introduces risks and responsibilities. With proper training, professionals can achieve far greater velocity, yet this pace can strain both systems and people. To unlock AI’s full potential and sustain performance, leaders and professionals must design thoughtful workflows, prioritize training, and build strong communities that support ethical and responsible adoption.

Chapter 7: Regulations and Governance

Canadian mortgage professionals face three overlapping layers of AI oversight:

Layer 1. Federal Lender Oversight – OSFI Guideline E-23

OSFI's final Guideline E-23, released September 11, 2025, takes effect May 1, 2027. The guideline establishes enterprise-wide model risk management requirements for all federally regulated financial institutions, explicitly including AI/ML systems within its scope.

Under E-23, lenders must implement risk-based governance frameworks that identify, assess, and manage model risks proportional to their impact. While the guideline uses broad language around "models," institutions have flexibility to determine which AI applications require full lifecycle governance based on their risk ratings and business materiality.

Although OSFI does not directly regulate mortgage brokerages, federally regulated lenders will need to demonstrate oversight of all models they use, acquire, or rely on—including those operated by third-party partners. This creates potential for contractual flow-down of governance requirements to brokerages, particularly for AI applications that materially impact lending decisions.

The extent of downstream requirements will depend on how individual lenders interpret and implement E-23's risk-based framework. Brokerages should expect requests for documentation around AI usage, risk management practices, and audit capabilities as lenders prepare for compliance.

Layer 2. Provincial Regulatory Frameworks

Provincial regulators are developing their own approaches to AI oversight in financial services. These frameworks typically focus on consumer protection, data privacy, and professional conduct standards. Requirements vary by province and continue to evolve as regulators assess the technology's impact on their sectors.

Layer 3. Federal Legislation Beyond OSFI

Beyond OSFI, existing federal laws already govern AI use: privacy legislation (Privacy Act, PIPEDA), FINTRAC's AML requirements, and Competition Bureau oversight. Additional AI-specific legislation remains under consideration.

Why Proactive Governance Matters

This multi-jurisdictional landscape makes AI governance a strategic business consideration rather than a simple compliance exercise. Brokerages that establish robust AI risk management practices early will be better positioned to meet evolving lender requirements, maintain funding relationships, and demonstrate their commitment to responsible innovation in an increasingly AI-driven industry.

Chapter 8: The Do's & Don'ts, Responsible Use and Guardrails

This section provides a practical framework for how to use Generative AI safely and effectively within DLCC Mortgage Group. The goal is to establish simple, habit-forming practices, creating a low-risk runway to build future-ready workflows that align with evolving governance expectations.

The Do's and Don'ts

TOP 3 DO's

- ✓ **Mask personal or sensitive data**
Never enter names, SINS, addresses, or deal-specific info. Always anonymize personal or sensitive information.
- ✓ **Validate AI outputs before you use them**
AI is a copilot, not autopilot. You remain the decision-maker.
- ✓ **Keep your AI work separate from personal accounts**
Always use a work-only account for business tasks.

Additional DO's

- ✓ Tag deal-related prompts with a simple file reference
- ✓ Save meaningful outputs (emails, summaries) in the deal folder/CRM
- ✓ Use fictional data when practicing or experimenting

TOP 3 DON'Ts

- ✗ **Don't upload borrower files or confidential data to public tools**
- ✗ **Don't trust AI outputs blindly and always validate before sending**
- ✗ **Don't mix personal and work AI accounts**

Additional DON'Ts

- ✗ Don't let AI draft final legal or compliance documents
- ✗ Don't use free consumer accounts for live lender-facing work
- ✗ Don't prompt anything you wouldn't show a lender auditor

Core Principles for Responsible Use of AI Models

- Mask personal or client-identifiable information before entering it into any AI model.
- Treat AI as a copilot, not autopilot. Always review, refine, and verify outputs before using them in client files or lender communications.
- Tag your prompts with deal IDs or file references. This small habit will make future logging seamless.
- Save meaningful outputs (e.g., client email drafts, underwriting summaries) in your deal folder.

Build Good Habits That Future-Proof Your Business

Banks may request documentation of AI usage as they implement OSFI E-23 requirements. By using AI models properly now, you will:

- Build prompt discipline
- Learn safe data handling habits
- Normalize lightweight prompt logging
- Reduce future compliance friction when lenders inevitably require full audit trails

Quick Reference: AI Usage Checklist

Before you hit “submit” on a prompt, ask:

- ✓ Am I using my work-only AI account?
- ✓ Did I mask or anonymize client information?
- ✓ Will I log this prompt/output in a way I can retrieve later?
- ✓ Would I be okay showing this exact prompt/output to a lender or regulator?

If you answer “no” to any, stop and adjust.

Thriving in the AI Era

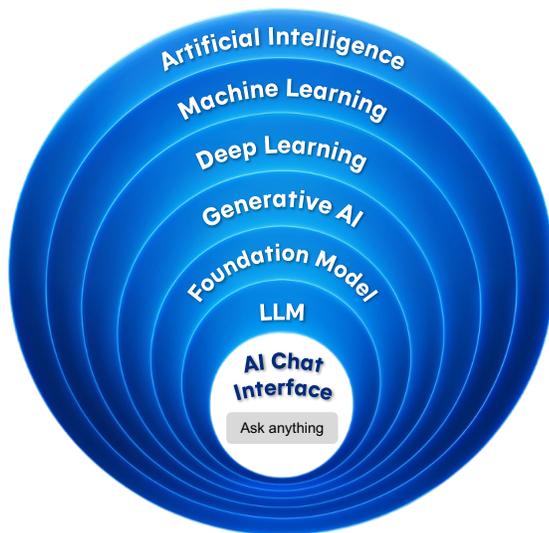
Learning to use generative AI competently and managing its regulatory realities go hand in hand. Those who adopt safe, auditable workflows now will protect lender trust and avoid potential challenges as governance requirements evolve.

This shift creates a clear opportunity. By moving early, building prompt discipline, masking data, and logging responsibly, you not only stay compliant but also gain a competitive edge. In a market where funding speed, reliability, and lender confidence drive success, mastering AI within the rules will separate top performers from the rest of the pack.

1. Human-AI Partnership

Generative AI is your **co-pilot**. It is not an autopilot or replacement. It accelerates routine work, freeing you to focus on clients, relationships, and strategic decisions. You stay in control. AI simply makes you faster, smarter, and more consistent.

2. Building Blocks



- **AI** – Computers performing tasks that normally require human intelligence
- **ML** – AI that learns patterns from data
- **DL** – Neural networks that process complex data
- **GenAI** – AI that creates new content (text, images, video, audio) from natural language
- **FM** – large versatile models trained on broad datasets and adapted for many tasks
- **LLMs** – Foundation Models specialized in text understanding and generation
- **Chat Interface** – The user-friendly way you interact with AI using natural language

3. Foundation Models for Work



4. What You Can Do With AI

RGI is the first principles framework for human-AI collaboration.

Retrieve: Ask | Find | Search

Generate: Write | Ideate | Create

Interact: Instruct | Query | Collaborate

Every meaningful **use case** can be understood through this lens.

5. How to Communicate With AI

If RGI defines what the AI can do then **Prompting** is the core communication skill that makes it work.

Prompting Formula:

Goal + Context + Format + Reference → **Refine**
 💡 Use pro techniques to improve output quality

CoT

Examples

Roles

Step-by-Step

Constraints

6. Benefits and Risks

| Benefits of Using AI | Risks of Using AI |
|-------------------------|---------------------------|
| ⬆️ Faster Workflows | ⚠️ Compliance / Privacy |
| ⬆️ Faster Learning | ⚠️ Hallucinations (false) |
| ⬆️ More Knowledge | ⚠️ Over-reliance on AI |
| ⬆️ Strong Relationships | ⚠️ Cognitive Overload |

7. The Do's & Don'ts

Do

- ✅ Mask personal or sensitive data
- ✅ Validate all AI outputs before use
- ✅ Use separate accounts for work

Don't

- ❌ Upload borrower files, names, SINs or personal information to public tools
- ❌ Don't trust AI outputs without review
- ❌ Don't mix personal and work accounts

IMPORTANT DISCLAIMER

This guide is intended for general informational purposes only. It provides a broad overview of considerations related to the use of generative AI in professional settings, including financial services. It is not tailored to the specific circumstances of any individual, firm, or jurisdiction and does not constitute legal, regulatory, or other professional advice. Readers should not rely on this guide as a substitute for personalized legal or compliance advice. Before implementing any generative AI tools or practices in your business, you are strongly encouraged to seek independent legal and regulatory advice relevant to your specific situation.