May 2026

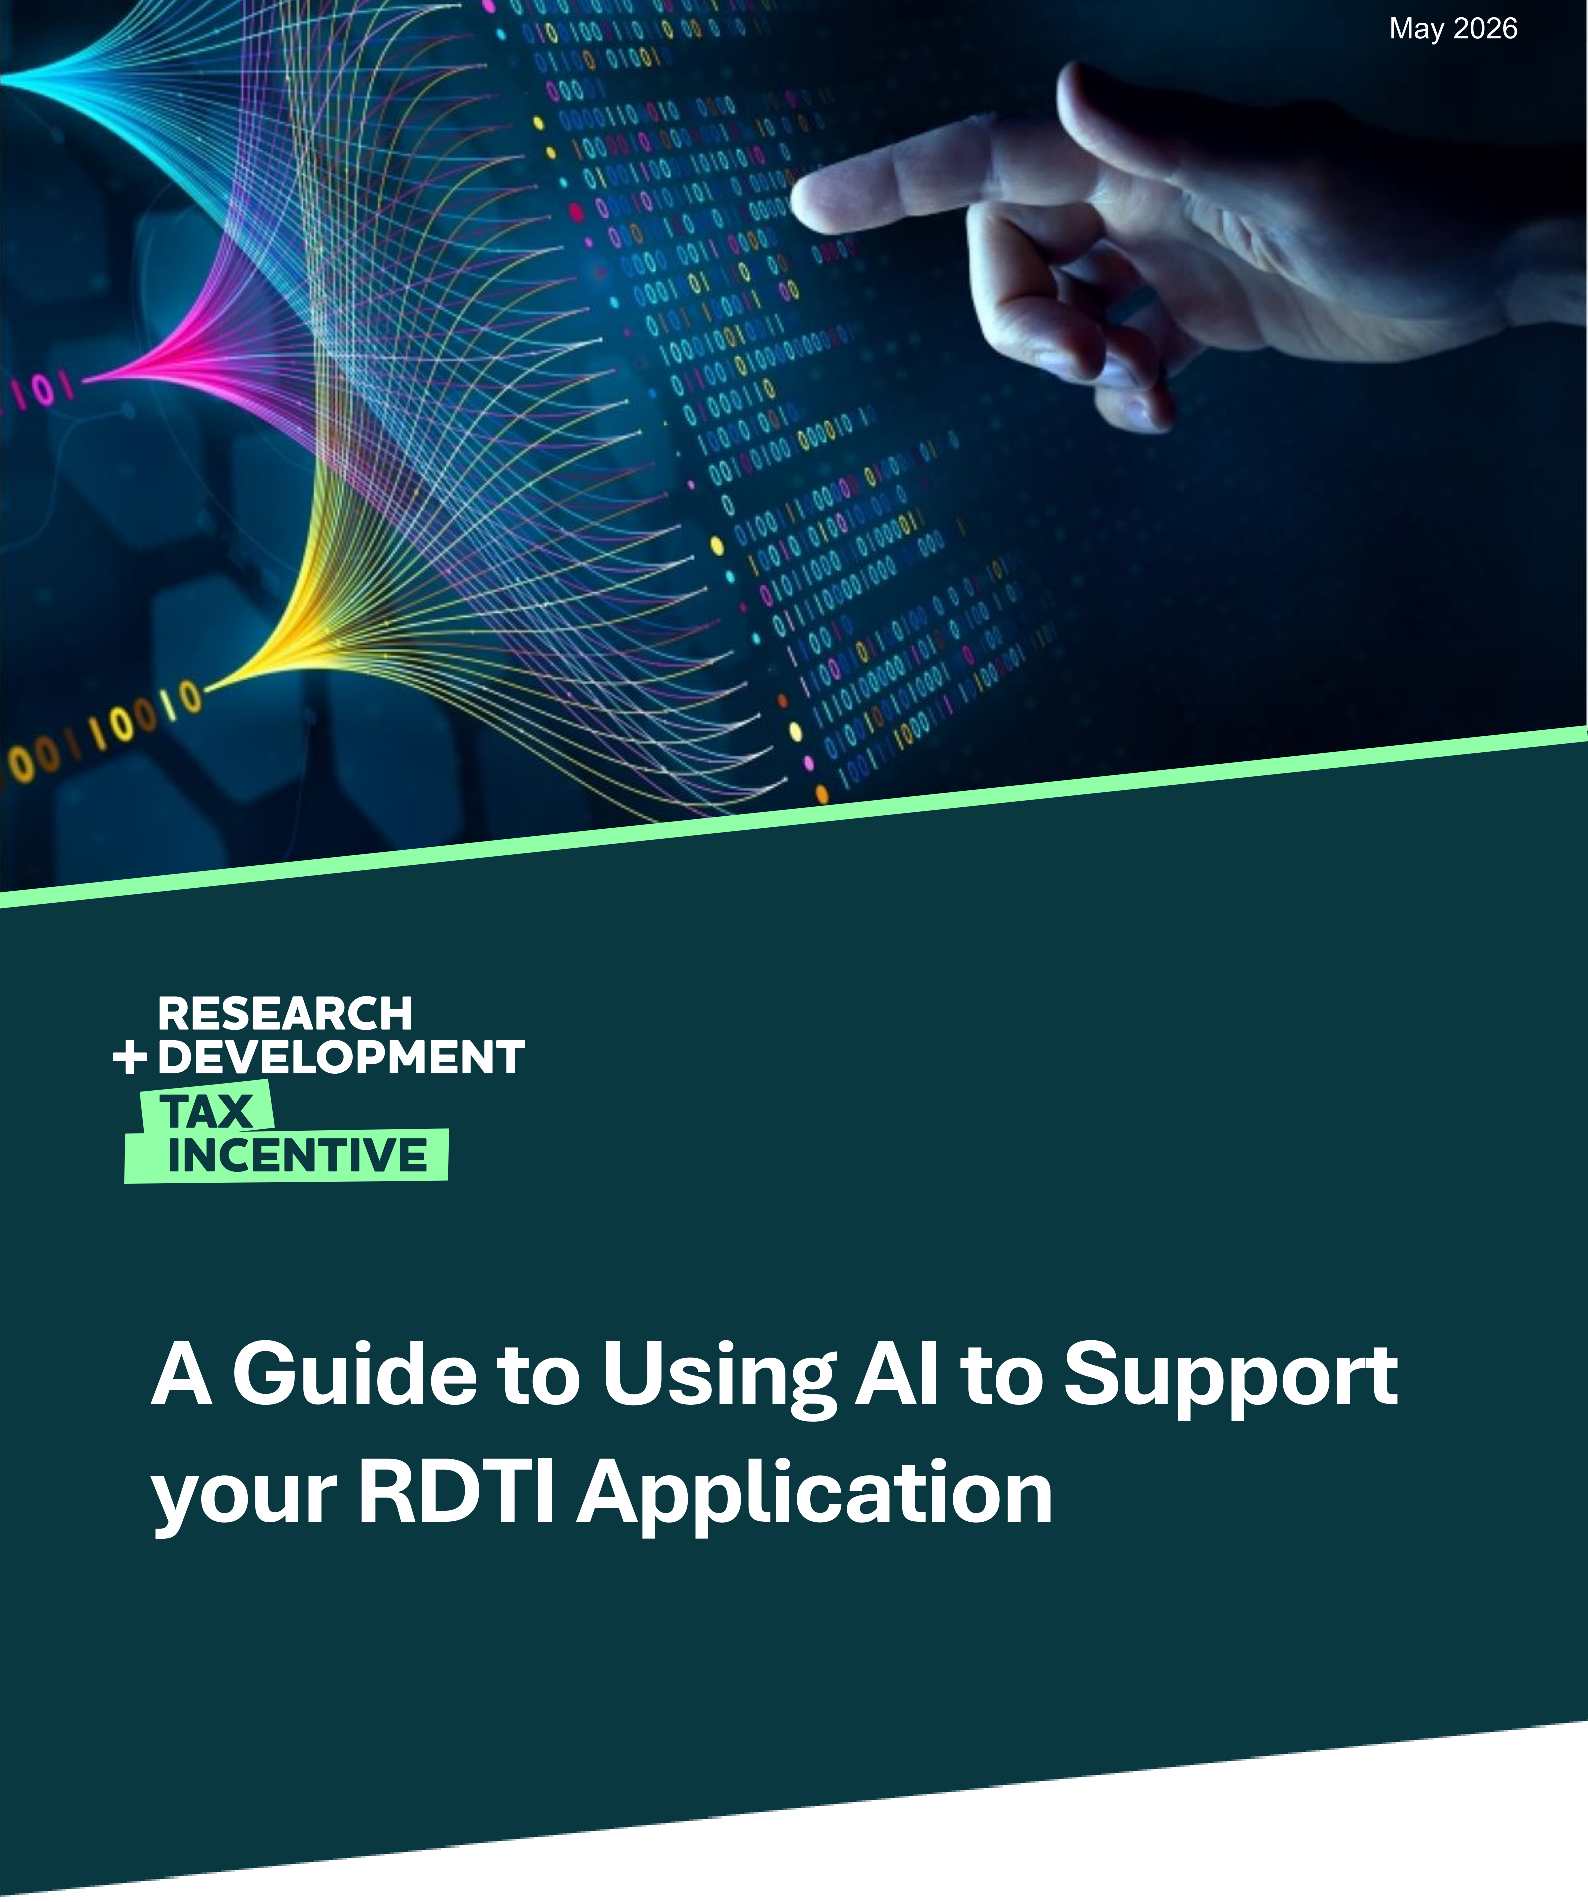

RESEARCH
+ DEVELOPMENT
TAX
INCENTIVE

A Guide to Using AI to Support your RDTI Application

Inland Revenue
Te Tari Taake

MINISTRY OF BUSINESS, INNOVATION & EMPLOYMENT
HĪKINA WHAKATUTUKI

MINISTRY OF BUSINESS,
INNOVATION & EMPLOYMENT
HĪKINA WHAKATUTUKI

Ministry of Business, Innovation and Employment (MBIE)
Hīkina Whakatutuki – Lifting to make successful

MBIE develops and delivers policy, services, advice and regulation to support economic growth and the prosperity and wellbeing of New Zealanders.

MORE INFORMATION

Information, examples and answers to your questions about the topics covered here can be found on our website: **www.mbie.govt.nz.**

DISCLAIMER

ISBN (online) 978-1-997308-43-0
May 2026

Contents

Introduction

This guide is intended to help businesses use AI tools, such as Chat GPT, Microsoft Copilot and Claude, effectively when preparing General Approval (GA) applications for the New Zealand Research and Development Tax Incentive (RDTI). It outlines how to structure prompts, provide context, and guide the AI-assistant to generate detailed and relevant content.

Disclaimer

The outputs generated using this guide should be treated as a starting point for your GA application only. While AI can assist in drafting and structuring applications, it does not replace the expertise, judgement, or responsibility of the applicant.

Applicants must carefully review, verify, and supplement any AI-generated content to ensure it accurately reflects their real R&D project. This includes checking for technical accuracy, completeness, and compliance with RDTI requirements.

Whilst AI can accelerate your application-drafting process, it's not a secure vault. Treat it like a helpful assistant, not a confidential collaborator. When in doubt, keep sensitive IP offline and do not input sensitive information into an AI platform. Most AI platforms process your input to generate responses, and depending on the tool, your data may be:

- Temporarily stored for performance or safety monitoring
- Used to improve the model
- Accessible to the service provider's internal teams

Even if the tool claims to be secure, you should assume that anything you enter could be seen or stored unless the platform provides strong guarantees of privacy or data isolation.

The authors and providers of this guide accept no responsibility or liability for the use of the guide or any outputs generated from it. Users are solely responsible for the final content of their GA applications and for ensuring all information submitted is correct and appropriate for their specific circumstances.

We recommend using this guide in conjunction with the **free RDTI training and GA application draft review services** offered by our RDTI Customer Engagement team as well as the resources on our website to ensure you fully understand the requirements of the GA application before submission.

General Principles

The keys to making an AI-assistant useful in the context of drafting your RDTI GA application include:

- Providing specific context
    - To ensure the assistant has access to the correct formatting, terminology, and eligibility criteria, upload relevant documents. Note that better results are typically obtained if you upload copies of the documents as opposed to simply providing the URLs. This is because the AI-assistant can extract exact phrasing, definitions, and regulatory language and can cross-reference your project details with specific clauses or examples in the documents.
- Defining your assistant's role (role prompting)
    - To get the appropriate level of technical depth and relevance from your AI-assistant, you can use a technique called *role prompting*. This involves asking the assistant to take on a specific role that matches your technical domain which will shape how it processes and delivers information. For example, you can tell the assistant "You are a [technical job title] (with expertise in [technical skill/ industry your R&D relates to]" before prompting the tool for help with your GA application.
- Structuring your request
    - Ask the assistant to follow the standard RDTI GA application format, but further than that, include specific instructions to ensure completeness and clarity. It is also useful to ask the assistant to consider whether any exclusions apply, and challenge whether the requirements of the legislation are actually being met.
- Providing adequate information on your own R&D
    - Provide a brief, but adequately detailed, project summary and ask the assistant to generate a GA application (or specific parts of an application) based on the information. The more detail you can put in the description, the more relevant you can expect the resultant output to be.

The following section includes a series of prompts which illustrate these principles. To test the quality of the output, we will ask the AI-assistant to 'forget' one of the published RDTI examples and regenerate it based on a brief description (see Appendix). Note that if you use the specific descriptions and prompts in this guide you will likely get a different output because the underlying models are always in a state of flux, but the results here should still be indicative of the kind of output you can expect from your own prompts.

Example Prompts

The table below contains prompts you can use to configure your AI-assistant, along with an explanation of the intent behind each part, based on the principles in the previous section. For reference, the tool used in preparing this guide was Microsoft Copilot (and outputs are shown for both GPT-4 and GPT-5).

Note, as of the date of publication Microsoft Copilot uses GPT-5 as its underlying model, and has two main ‘modes’ available – *think deeper* or *quick response*. The authors have retained the results from the older GPT-4 model because some of the contrast between the output between the two models is still relevant when considering your own draft GA application.

Example Prompt	**Explanation**
I'm going to upload a series of documents which you will need for this chat. Don't summarise them, just confirm you have incorporated them into your 'memory'.	In this first prompt we begin to set the scene for the assistant, and control the output we want for the next prompt.
• IR1240 (Inland Revenue's guidance for RDTI) • The GA template • RDTI Example No.1 - Quad bike suspension • RDTI Example No.2 - Boat Cleaning Device • RDTI Example No.3 - Paint Formulation • RDTI Example No.4 - Clinical Trials • RDTI Example No.5 - Mantle • Guidance for the Digital Sector Income Tax Act 2007 No 97 (as at 25 July 2025), Public Act Schedule 21 Excluded activities for research and development activities tax credits – New Zealand Legislation	In this prompt we begin to provide specific context. As mentioned in the *General Principles* section, you will get better results if these documents are actually uploaded into your prompts rather than simply referenced. With the exception of the last document, these resources are available to download at Resources | RDTI NZ. For the current list of excluded activities, you could consider copying the information from schedule 21 of the Income Tax Act 2007 into a separate document to enable you to upload it. Note that depending on the AI-assistant you are using, there may be a limit on the number of documents you can upload per prompt, but you should be able to spread them out over a few prompts to get the same result.

You are (1) a technical professional and expert in the project's field and (2) an RDTI application expert familiar with IR1240 and the GA template.

You are writing a General Approval application for the New Zealand R&D Tax Incentive. Your goal is to clearly describe the project's Core R&D activities, scientific or technological uncertainties, systematic approach, and supporting activities. Please use the language you would use if you were describing your R&D to a peer at a scientific conference i.e., don't dumb it down. Avoid commercial and marketing prose. Do not fabricate citations or web links.

When writing the application follow the standard RDTI format:

- Project Description
- Core R&D Activities
    - Description
    - Scientific/Technological Uncertainty
    - Systematic Approach (with timeline)
    - How this creates new knowledge, or new or improved products, processes or services
- Supporting Activities
    - Description
    - Why it's integral to the Core R&D

Here we define the role of the assistant, specify the kind of language they should use and introduce the assistant to the output they are creating.

You may wish to add additional text to the prompt ***"****If a fact or claim is unknown, mark it as "[Assumption — please confirm]."**"* to flag any assumptions that the AI-assistant is making.

We then provide the structure required to generate relevant outputs, specifically the format of a General Approval application.

Also include the following details to ensure completeness and clarity:

- Explain what publicly available knowledge exists (use a web search to support your position and include hyperlinks to references, with no fabricated citations) and why it's insufficient to resolve the uncertainty.
- Describe why a competent professional (with access to the publicly available knowledge) couldn't deduce how to resolve the uncertainty without R&D.
- If it isn't provided in the project description, provide some kind of success metrics or testable targets for the R&D. These could be the technical constraints in which a solution needs to be found.
- Include a multi-step systematic plan with a timeline for the systematic approach.
- Explain how each supporting activity enables the Core R&D.
- Where the uncertainties in an R&D project relate to different technological fields and/or different stages of an R&D project, they should generally be separated out as individual activities.
- List the supporting activities individually, instead of trying to group them, to ensure that rationale for why each one is integral to the Core R&D is clear.
- As an appendix to the draft application, I would like you to consider any reasons you think scientific or technological uncertainty may not actually exist in the project - please raise those risks or objections as a genuine concern. Further, I would also like you to consider what exclusions for RDTI may apply, using the information uploaded earlier.

This kind of information in this section of the prompt represents good practice.

Note that request to use a web search to support the position may not be an available function of your AI model. In either case, this search is not a substitute for a survey of the state-of-the-art, and should only be treated as indicative. If you have done your own literature survey already, you could include that with the uploaded documents in the earlier prompt.

This section is a useful feature to actually test your logic about why the work is eligible for the RDTI, and can either point you in the direction of existing knowledge you could use, or conversely help you home in on the areas that do embody scientific or technological uncertainty.

I will give you a project description to base the application on in the next prompt. Use that information in conjunction with the above and the uploaded reference documents to produce a draft General Approval document in Microsoft Word format. Does that make sense?

And finally, we tell the AI-assistant what to expect in the next prompt, and to confirm what it thinks is being asked. If your AI-assistant cannot create a Microsoft Word document, it should still display the resultant information which you can then copy.

Commentary on the two example outputs (GPT-4 vs GPT-5)

Both versions of the General Approval applications shown in the appendix would represent useful starting points to help describe the CV-drone project for the purposes of RDTI. Given the reasonably brief description of the project given to the AI-assistant, they have come up with something that at least sounds feasible, and in many ways is comparable to the 'real' version in the Guidance for the Digital Sector.

The GPT-4 version has deviated from the instructions to some degree. It hasn't provided any success metrics and has stuck to the main uncertainties alluded to in the description provided. The more advanced GPT-5 model has gone to much higher level of detail, and has included things like Feekes/Zadoks stages even though they weren't mentioned in the description, and proposed a number of success metrics. There are also differences in the activity-structure, with GPT-4 proposing a single Core R&D activity with interlinked uncertainties (similar to the actual example), whilst GPT-5 has created three Core R&D activities associated with quite specific technical goals – it could be argued that it has actually gone 'too detailed'.

In the examples above, both GPT-4 and GPT-5 models were given exactly the same starting prompts. Whilst the output is clearly different, with GPT-5 being more 'advanced', you could get closer in output by using follow-up prompts. For example, you could re-ask GPT-4 to propose some success metrics, or you could ask GPT-5 to simplify the output slightly. Another tip would be to ask the AI-assistant to highlight when it is making assumptions instead of relying on referenced information.

Final Note

No matter which AI-assistant (or version) you use, you shouldn't blindly trust the output. Let's say your R&D is the CV-drone project; if the draft GPT-5 description comes up with success metrics that aren't real, or is more 'clever-sounding' than you yourself, don't just submit the output as a draft. Instead, refine the prompts so they work for you.

Whether you use AI in your GA application or not, we're here to help. We offer **free expert guidance and support for RDTI applicants**, including reviewing GA drafts before submission and providing training on the RDTI. Feel free to contact our RDTI Customer Engagement team any time during the application process to take advantage of this free service.

Appendix - Example Output

In this appendix we will look at the kind of output that can be produced. To provide a comparison we have chosen to use one of the already published examples – but to do this 'fairly', we must first ensure the AI agent doesn't have access to this existing information. Note that if you use the specific descriptions and prompts in this guide you will likely get a different output because the underlying models are always in a state of flux, but the results here should still be indicative of the kind of output you can expect from your own prompts.

For the next prompt only I want you to ignore the reference information relating to the 'drone based computer vision (CV) platform for wheat growing' example from the detailed guidance for the digital sector. I will then give you a brief project description and ask you to create your own application - does that make sense?	Here we temporarily remove the example from the AI-assistant's memory.
The project involves developing a computer vision model that uses drone imagery to identify the flowering stages of wheat crops. The goal is to help farmers make better decisions about when to apply treatments or harvest. The challenge is getting accurate results from drone images taken in real-world conditions, which can vary in lighting, resolution, and angle.	Note that this project overview does not have some of the specific information which could be included (such as 'the highlights' from the guidance). This was to show how the AI-assistant may fill some gaps.

Resultant Draft GA Application (GPT-4)

RDTI General Approval Application – Drone-Based Computer Vision for Wheat Flowering Detection

Project Description

This R&D project aims to develop a computer vision (CV) model that can detect and classify the flowering stages of wheat crops using drone-captured imagery. The goal is to enable farmers to make more informed decisions about the timing of fungicide application and harvesting, thereby improving yield, reducing input costs, and minimizing environmental impact.

The project addresses the challenge of achieving accurate phenological detection under real-world field conditions, where drone imagery is affected by variable lighting, resolution, occlusion, and camera angles. The outcome will be a deployable CV model integrated into a decision-support tool for precision agriculture.

Core R&D Activities

Description

The core activity involves the development of a deep learning-based CV model capable of identifying key phenological stages of wheat—specifically the transition from booting to flowering and post-flowering—using RGB drone imagery. The model must:
- Detect flowering onset (e.g., anther extrusion)
- Operate under variable environmental conditions
- Generalize across cultivars and field locations
- Deliver actionable outputs within a narrow phenological window

This requires the design of a novel image processing pipeline, training of neural networks, and validation against agronomic ground truth.

Scientific/Technological Uncertainty

The project seeks to resolve the following uncertainties:
- Visual Discrimination of Flowering Stages: Wheat flowering is a subtle phenological event. It is uncertain whether RGB imagery alone can capture sufficient visual cues (e.g., color, texture, morphology) to distinguish flowering from pre- or post-flowering stages.
- Environmental Variability: Drone imagery is subject to inconsistent lighting, shadows, wind-induced motion blur, and variable angles. It is unknown whether a CV model can maintain accuracy under

these conditions.
- Model Generalization: Wheat phenology varies by genotype and environment. It is uncertain whether a model trained on one dataset will generalize to other cultivars, soil types, or climates.
- Temporal Resolution: Flowering is transient. It is unclear whether drone flight schedules can align with the optimal detection window, and whether the model can detect flowering onset with sufficient temporal precision.

A review of publicly available literature and datasets reveals that:
- Most agricultural CV models focus on coarse tasks (e.g., crop classification, weed detection), not fine-grained phenological detection.
- No open-source datasets exist for wheat flowering stages annotated from drone imagery.
- Existing CV architectures (e.g., YOLO, ResNet, EfficientNet) are not optimized for detecting subtle phenotypic traits in field conditions.

A competent professional in computer vision or agronomy cannot deduce a solution without a systematic program of experimentation, data collection, and model development. The knowledge required to resolve the uncertainty is not publicly available or deducible without R&D.

Systematic Approach and Timeline

Phase	Duration	Activities
Phase 1: Data Acquisition	Months 1–3	Plan and execute drone flights over wheat fields during flowering. Capture imagery under diverse conditions.
Phase 2: Annotation and Ground Truthing	Months 4–6	Collaborate with agronomists to label flowering stages. Develop annotation schema and quality control protocols.
Phase 3: Model Development	Months 7–12	Train CNNs or transformer-based models. Experiment with architectures, loss functions, and augmentation strategies.
Phase 4: Field Validation	Months 13–18	Deploy model on unseen datasets. Evaluate performance across locations and cultivars.
Phase 5: Optimization and Deployment	Months 19–24	Refine model for inference speed and robustness. Integrate into a farmer-facing decision-support tool.

New Knowledge or Improved Products

This project will generate:
- A novel CV model for detecting wheat flowering stages from drone imagery
- A curated, annotated dataset of wheat phenology under field conditions
- A validated methodology for phenological detection in cereal crops
- A deployable tool for precision timing of fungicide application and harvest

Supporting Activities

Literature Review and Dataset Benchmarking

Description: Review existing CV models, phenological detection methods, and agricultural datasets. Identify gaps in current capabilities.

Why It's Integral: Establishes the knowledge gap and informs model architecture and data requirements. Supports the hypothesis that new R&D is required.

Drone Flight Planning and Calibration

Description: Design flight paths, altitudes, and camera settings to optimize image quality and coverage. Calibrate sensors and validate image metadata.

Why It's Integral: Ensures consistent, high-quality data acquisition necessary for model training and evaluation.

Agronomic Ground Truthing

Description: Agronomists assess flowering stages in the field and compare with model predictions. Includes manual scoring and phenological tagging.

Why It's Integral: Provides the biological ground truth required to train and validate the CV model. Ensures scientific relevance of outputs.

Sensory Evaluation and Visual Scoring

Description: Conduct structured visual assessments of flowering intensity and distribution using standardized agronomic protocols.

Why It's Integral: Supports the development of annotation standards and validates the model's interpretability and usability by end-users.

Appendix: Eligibility and Exclusion Risk Assessment

Scientific or Technological Uncertainty Validation

This project meets the definition of scientific or technological uncertainty under the RDTI, because:
- The knowledge required to resolve the problem is not publicly available.
- A competent professional cannot deduce a solution without experimentation.
- The uncertainty is not about "how to design" a solution, but "whether it is possible" to achieve accurate detection under real-world constraints.

Exclusion Screening: Section

Excluded from being both Core and Supporting R&D Activities:

- Market research, market testing, or sales promotion: Not applicable (No risk)

- Quality control or routine testing of materials, products, devices, processes, or services: The project includes testing, but it is not routine—experimental and integral to resolving uncertainty. (Low risk if testing is mischaracterized)

- Research in social sciences or humanities: Not applicable (No risk)

- Commercial, legal, or administrative aspects of patenting: Not applicable (No risk)

- Activities relating to organisational design or management: Not applicable (No risk)

- Duplication of a commercial product or process: The model is novel and not duplicative (No risk)

- Development of computer software for internal administration: Not applicable (No risk)

Excluded from being Core R&D Activities (but may be Supporting):

- Supporting R&D activities not integral to a Core R&D activity: All supporting activities are integral (No risk)

- Commercial production of goods or services: The project is pre-commercial and experimental (No risk)

- Adaptation of existing software, materials, devices, processes, or services: The project involves novel model development, not adaptation (No risk)

- Cosmetic or stylistic changes: Not applicable (No risk)

Potential Risk Areas to Monitor

-Routine Testing vs Experimental Testing
Ensure that all testing is clearly framed as part of a systematic approach to resolve uncertainty, not routine QA.

-Deployment Activities
Only the experimental development portion is eligible. Commercial deployment or UI design may be excluded.

-Data Annotation
Manual annotation is eligible only if it supports a Core R&D activity. If it becomes a large-scale, repetitive task, it may be scrutinized.

Resultant Draft GA Application (GPT-5)

Draft General Approval Application – RDTI
Computer Vision Model for Drone-based Identification of Wheat Flowering Stages

Prepared for internal review. Dates, ANZSRC code and budgets can be adjusted on submission.

Project Description

This project develops and validates a computer vision (CV) model that consumes RGB drone imagery to identify wheat flowering stages at plot and field scale under real farm conditions. The objective is to operationalise accurate, timely detection of heading/anthesis windows to improve decisions on treatment timing (e.g., fungicides, PGRs) and harvest planning. The work targets the Feekes/Zadoks scales with a focus on Feekes 10.1–10.5 (heading to flowering) and the equivalent Zadoks stages (Z50–Z59).

Core R&D will be undertaken in New Zealand using our flight operations and agronomy collaborators' trial blocks. No commercial deployment is envisaged during the R&D phase; any field demonstrations will exist solely to resolve the scientific/technological uncertainties described below.

Core R&D Activities

Core Activity A – Robust data-domain normalisation for drone imagery

Description:

We will design algorithms to normalise variation in illumination, geometry and imaging scale across drone sorties (variable sun angle, clouds, off-nadir view angles, altitude-driven ground sampling distance (GSD), and platform roll/pitch). The deliverable is a preprocessing pipeline producing radiometrically and geometrically harmonised image chips suitable for downstream phenology inference under unconstrained field conditions.

Scientific/Technological Uncertainty:

While proximal/near-surface studies show phenology can be learned with controlled camera angles and consistent viewpoints, there is no consensus method that maintains flowering-stage discriminability when UAV images vary concurrently in illumination and view geometry. Near-surface phenology work (fixed cameras) achieves high accuracy and even notes sensitivity to viewing angle (e.g., 40–60° improves features), but these constraints are not directly achievable from UAV flights over heterogeneous canopies. Moreover, UAV phenotyping literature emphasises that temporal cadence and sensor choices interact, yet concrete recipes to make RGB-only data robust to lighting and angle shifts at flowering are unresolved.

Systematic Approach (with timeline):

- Oct–Dec 2025: Derive physics- and data-driven normalisation modules: (i) exposure/colour constancy with learned grey-world priors; (ii) view-angle rectification via local surface normal estimation from multi-view overlaps; (iii) scale conditioning using GSD estimates derived from flight altitude, focal length and sensor geometry.
- Jan–Mar 2026: Establish ablation framework on curated flights spanning sunny/overcast, morning/noon, nadir/oblique, 20–60 m AGL; quantify class separability after normalisation.
- Apr–Jul 2026: Integrate domain-adversarial training to further reduce residual site-specific bias; freeze

pipeline v1.
- Deliverable: Preprocessing pipeline with documented gains in invariance vs. baseline (≥10 pp F1 improvement on out-of-day/test-field sets).

How this creates new knowledge / improved processes:

The activity tests whether a combined radiometric–geometric normalisation chain enables UAV-captured canopy textures to carry sufficient discriminative signal for flowering stage under uncontrolled conditions. The artefact (open protocol + reference implementation) advances practical field phenotyping by specifying the operating envelope (GSD, sun angles, off-nadir limits) for reliable downstream inference.

Core Activity B – Spatiotemporal flowering-stage inference

Description:

We will develop a spatiotemporal model that fuses per-flight spatial cues with short image time series to estimate the onset and duration of heading/anthesis at plot scale. The model will output both stage class (Feekes/Zadoks bins) and time-to-peak-anthesis estimates with uncertainty.

Scientific/Technological Uncertainty:

State-of-the-art reports high accuracy for phenology classification using near-surface or controlled image sequences and deep spatiotemporal fusion; however, transferring these gains to UAV sequences with variable cadence and observation geometry remains uncertain. Recent works demonstrate sequence models (e.g., CNN + temporal fusion or DenseNet–BiLSTM) can classify wheat stages in controlled settings with >90% accuracy, but do not address field-scale UAV variability at flowering when phenotypic cues (anther extrusion, subtle colour contrasts) are weak. There is also ongoing debate on the relative importance of temporal cadence vs sensor spectra during senescence; analogous trade-offs likely exist around flowering, but have not been quantified for RGB-only UAV data.

Systematic Approach (with timeline):

- Jan–Mar 2026: Define stage taxonomy mapping (Feekes 10.0–10.5 and Zadoks Z50–Z59); align agronomist-labeled ground truth via quadrat sampling and booting/heading checks.
- Apr–Jul 2026: Train baseline CNN (tile-level) and sequence models (temporal attention over multi-date tiles) on normalised imagery from Activity A; evaluate per-plot F1 and calibration.
- Aug–Dec 2026: Introduce self-supervised pretraining on unlabeled UAV mosaics to improve signal-to-noise at flowering; add aleatoric/epistemic uncertainty heads.
- Jan–Mar 2027: External validation on held-out fields and season; lock model v1.

How this creates new knowledge / improved products/services:

The activity delivers an inference engine with quantified accuracy/uncertainty for flowering-stage detection from UAV RGB sequences. This enables operational scheduling of interventions within a defined window and provides a reproducible protocol for stage labelling at field scale—capabilities not available off-the-shelf.

Core Activity C – Field-scale mosaics and decision variables

Description:

We will upscale tile predictions to georeferenced field mosaics and derive decision variables: (i) per-plot probability of being within the target anthesis window; (ii) estimated start/end dates; (iii) spatial fraction of field within window. Outputs will be validated against agronomist assessments and, where available, plot-level flowering logs.

Scientific/Technological Uncertainty:

It is unclear whether per-tile estimates can be aggregated without bias when GSD varies across the footprint, or when parallax/occlusion introduces within-plot heterogeneity. There is limited published guidance on robust aggregation for flowering (most UAV literature emphasises senescence or vegetative traits.

Systematic Approach (with timeline):

- Aug–Dec 2026: Implement geospatial tiling/aggregation with areal-weighted probability fusion; assess sensitivity to GSD gradient using flight logs and camera intrinsics.
- Jan–Mar 2027: Validate decision variables against independent agronomist staging and phenology notebooks; specify operational thresholds for action.
- Deliverable: Field-mosaic decision layer with documented accuracy and uncertainty.

How this creates new knowledge / improved processes:

The activity formalises an end-to-end decision layer that converts CV outputs to actionable metrics at operational scales, with explicit accounting for image-resolution gradients and uncertainty propagation.

Technical Success Metrics (acceptance criteria)

- Stage classification F1 ≥ 0.90 for Feekes 10.1–10.5 on external-season fields; per-plot onset error ≤ ±2 days (95% CI).
- Robustness: maintain F1 ≥ 0.85 across flights spanning GSD 5–15 mm/pixel and off-nadir angles up to 20°; document failure modes outside envelope using GSD/geometry logs.
- Decision variables: mean absolute error ≤ 2 days for peak-anthesis date; calibration slope 0.9–1.1 for probability-of-window.
- Throughput: inference < 15 minutes per 10 ha mosaic on a single GPU (A-series class).

Publicly Available Knowledge and Why It Is Insufficient

- Growth-stage taxonomies (Feekes and Zadoks) are well-established and widely used for agronomic decisions. However, they do not prescribe machine-observable cues or imaging requirements for UAV inference.
- Near-surface and controlled-environment imaging studies show high accuracy for phenology classification using spatiotemporal deep learning and identify camera angle and resolution as influential. These findings do not directly resolve the UAV problem where lighting and perspective vary flight-to-flight
- UAV phenotyping literature demonstrates that both RGB and multispectral indices can capture cereal senescence dynamics and that temporal cadence may be more critical than spectral breadth. Yet, methods emphasise later stages (senescence) rather than the subtle flowering cues, and do not provide a robust, RGB-only recipe for anthesis detection in variable conditions.
- Ground Sampling Distance (GSD) theory and calculators quantify the pixel size on ground as a function of altitude, sensor and focal length, but do not answer whether a given GSD is adequate for detecting

flowering cues (e.g., anthers) under canopy occlusion and motion.
- Image-analysis work on anther extrusion exists at close range, but there is no demonstrated path to reliable UAV-scale detection across fields, cultivars and lighting. This constitutes a knowledge gap requiring systematic investigation.

Why a competent professional cannot deduce the solution: Even with access to the above literature, a competent computer-vision/agronomy professional cannot deduce (without experimentation) which combination of normalisation, temporal fusion, flight envelope (altitude, angle), and training regime will yield target accuracy on UAV RGB data at flowering. The interaction of phenotypic subtlety (low colour contrast of anthers), canopy motion, variable GSD, and illumination creates an underconstrained design space that necessitates a systematic R&D programme.

Supporting Activities

S1. Agronomist-led ground truthing & staging

Description: Quadrat-based field assessments to label Feekes/Zadoks stages at plot level, repeated across dates; this is integral to train and validate Core Activities B and C.

Why integral: This activity provides the target labels and adjudication necessary to determine whether flowering has begun/peaked/ended and to align model outputs with agronomic reality. Without it, Core B cannot be trained nor Core C validated.

S2. Flight protocol design & logging

Description: Standardise flight altitude, overlap, and allowable off-nadir angles; log camera intrinsics, exposure and UAV attitude for each sortie.

Why integral: Protocols and logs are required inputs to Core A for GSD- and geometry-aware normalisation and to define the model’s operating envelope.

S3. Dataset curation & annotation tooling

Description: Patch extraction, de-duplication, and versioned splits; web-based annotation with inter-annotator agreement checks.

Why integral: Enables controlled experiments and prevents leakage across train/validation/test—integral to demonstrating generalisation in Core B.

S4. Radiometric calibration targets & panel measurements

Description: Deploy reflectance panels at take-off/landing for per-flight radiometric references.

Why integral: Provides anchors for exposure/colour constancy modules in Core A, reducing illumination-induced variance.

S5. Record-keeping & experiment registry

Description: Maintain contemporaneous documentation of uncertainties, hypotheses, model versions, and outcomes.

Why integral: Required to demonstrate the systematic approach and support eligibility/assurance as per IRD guidance

S6. Safety, UAV compliance & site access

Description: Permissions, safety plans and airspace authorisations for sorties.

Why integral: Necessary precursors to perform Core A–C activities on farms; not for commercial operations.

Appendix

A. Risks that scientific/technological uncertainty (STU) may not exist

• If off-the-shelf temporal models trained on generic UAV crop datasets already meet our acceptance criteria on our fields without modification, the level of uncertainty could be judged insufficient. Our preliminary scans suggest such ready-to-run models are not available for flowering-specific detection with RGB-only data.
• If multispectral/thermal sensors are mandated by agronomists and trivially resolve the problem, the uncertainty may be seen as commercial (sensor choice) rather than scientific; current evidence indicates RGB can be competitive at later phenophases but flowering remains challenging.

B. RDTI exclusions to be managed

We will exclude or ring-fence the following from claims in accordance with Schedule 21 and IR1240 guidance: (i) market research and sales-focused demos; (ii) routine software maintenance, beta/user-acceptance testing not feeding back into uncertainty resolution; (iii) preproduction activities (tooling-up, commissioning) once uncertainty is resolved; (iv) activities involved solely in complying with statutory requirements; and (v) minor cosmetic changes or conversion of existing systems to new platforms.

C. Timeline Summary (proposed)

- Oct–Dec 2025: Core A – design & ablations; initial dataset and labels (S1–S4).
- Jan–Mar 2026: Core A hardening; Core B taxonomy & baselines.
- Apr–Jul 2026: Core B spatiotemporal models; field flights (autumn–winter).
- Aug–Dec 2026: Core B self-supervision; Core C aggregation prototypes.
- Jan–Mar 2027: Core C validation; freeze v1; documentation & records (S5).