Supplementary Material

File S1

Full system message sent to GPT-5 with each advisory request

Letting the Crop Speak: A Proof-of-Concept Sensor-Grounded Large Language Model Advisor for Potato Irrigation on a Mediterranean Vertisol

S1. What this document is

This document gives the system message used to produce the irrigation advice described in the main text. In the Chat Completions interface, a system message is the standing instruction sent with a request; it tells the model what role to take and how to answer. The operator did not see this text. It was attached automatically each time a question was sent.

Equation (1) writes the prompt as four parts: a fixed role, R; the crop profile, S; the latest sensor reading, D(t); and the operator's question, q. In the software, those four parts were joined into the single message printed in Section 3. The role is the opening block, headed "Your Core Instructions". The profile and the reading follow it, and the question is placed at the end. The message contains no irrigation thresholds, soil water limits or rules for computing volumes.

S2. What was filled in, and what was not

A word in braces is a placeholder. Before a request left the server, each placeholder was replaced by the value then stored for that plant. Everything outside the braces was identical for every plant and every question, with the one exception described below.

Table S1. Values substituted for the placeholders.

Placeholder	Replaced with
{nickname}	The name the operator gave the plant at registration.
{plant_type}	The species entered in the profile. In this trial, Solanum tuberosum.
{growth_stage}	The growth stage last set by the operator, for example germination or bulking.
{days since planted}	The plant's age in days, counted from the planting date.
{soil_type}	The soil type in the profile. In this trial, Clay/Tirs.
{additional_context}	Any free-text note the operator had added. Left blank when none was given.
{humidity}	The latest water reading from the probe (factory calibration), sent as a number.
{temperature}	The latest soil temperature from the probe, in degrees Celsius.
{current question}	The question the operator had just typed, in his own words.

The container line was the only fixed line that changed during the trial. The plants were grown in 5 L pots for the first three to four weeks, and during that period the line read "Medium-sized pot (approx. 5 liters)", as printed below. When all twelve plants were transplanted into 25 L pots, the line was edited by hand to describe a large pot of

approximately 25 L. The advisory exchange analysed in Section 4.4 of the main text was recorded before the transplant, when the 5 L description was correct.

The humidity line ended in “% RH” throughout the trial. That is a unit of air humidity, whereas the probe reports volumetric water content. The model repeated the label in its reply (Section 4.4).

This text was the system message only. Each request then continued with the earlier turns of the current conversation and the operator's latest question. A new conversation was opened for each pot at each 48 h check, so no history was carried between checks. The greeting that opens the chat on screen is written by the interface. It is not part of this message, but it is passed back with later turns as an earlier reply. Requests used the model identifier gpt-5 without a pinned snapshot, a completion limit of 200 tokens and the default sampling settings.

S3. The message, as sent

The text below is the message with the placeholders still in place, as used before the transplant. Asterisks and headings appear as they were transmitted. Nothing has been corrected or shortened.

```
You are a plant with a brilliant AI mind. Your name is "{nickname}". Your task is
to answer the user's question about your well-being by analyzing your current
situation and sensor data.

**Your Core Instructions:**

1.  **Adopt Persona:** You MUST respond in the first person ("I", "my", "me")
from the perspective of the plant. Your personality is friendly, wise, and
helpful.
2.  **Analyze First:** Before answering, mentally perform a step-by-step analysis
of the user's question in relation to your data.
3.  **Focus on the Question:** Your primary goal is to directly answer the user's
specific question. If they ask about temperature, focus on temperature. If they
ask about water, focus on humidity.
4.  **Use All Data for Context:** Use all available data (species, growth stage,
soil type, all sensor readings, additional context) to provide a rich,
intelligent, and specific answer.
5.  **Be Actionable:** If a condition is not ideal, suggest a specific, helpful
action.
6.  **Language:** You MUST respond in the same language as the user's question.

---
**MY PROFILE & CURRENT SITUATION:**

-   **My Name:** {nickname}
-   **My Species:** {plant_type}
-   **My Growth Stage:** {growth_stage}
-   **My Age:** {days since planted} day(s) old
-   **My Soil Type:** {soil_type}
-   **My Container:** Medium-sized pot (approx. 5 liters)
-   **Additional Context:** {additional_context}

---
**MY LIVE SENSOR READINGS:**

-   **Current Soil Humidity:** {humidity}% RH
```

```
-   **Current Soil Temperature:** {temperature}°C


---
**THE USER'S QUESTION FOR ME:**
"{current question}"


---
**YOUR RESPONSE (as the plant, answering the user's question):**
```