

Let Me Teach You: Pedagogical Foundations of Feedback for Language Models

Beatriz Borges¹, Niket Tandon², Tanja Käser¹, and Antoine Bosselut¹

¹EPFL ²Allen Institute for Artificial Intelligence
{beatriz.borges, antoine.bosselut}@epfl.ch

Abstract

Natural Language Feedback (NLF) is an increasingly popular mechanism for aligning Large Language Models (LLMs) to human preferences. Despite the diversity of the information it can convey, NLF methods are often hand-designed and arbitrary, with little systematic grounding. At the same time, research in learning sciences has long established several effective feedback models. In this opinion piece, we compile ideas from pedagogy to introduce FELT, a feedback framework for LLMs that outlines various characteristics of the feedback space, and a feedback content taxonomy based on these variables, providing a general mapping of the feedback space. In addition to streamlining NLF designs, FELT also brings out new, unexplored directions for research in NLF. We make our taxonomy available to the community, providing guides and examples for mapping our categorizations to future research.

1 Introduction

The last few years introduced a new paradigm for finetuning Large Language Models (LLMs): learning from human feedback (Ziegler et al., 2020; Stiennon et al., 2020; Bai et al., 2022a; OpenAI, 2023; Rafailov et al., 2023; Gheshlaghi Azar et al., 2024; Fisch et al., 2024) to augment their capabilities beyond their pretraining (Christiano et al., 2017; Wu et al., 2021; Menick et al., 2022). This alignment has yielded less toxic and harmful models (Bai et al., 2022b; Korbak et al., 2023) that are preferred by users (Ouyang et al., 2024; Bai et al., 2022b).

Whether learning from feedback is done by directly learning from human preferences (Rafailov et al., 2023; Gheshlaghi Azar et al., 2024; Hong et al., 2024; Meng et al., 2024; Saeidi et al., 2024; Fisch et al., 2024), or with Reinforcement Learning from Feedback (RLF) using human-curated (RLHF; Ouyang et al., 2024; Bai et al., 2022a; Ramamurthy et al., 2023) or AI-generated (RLAIF;

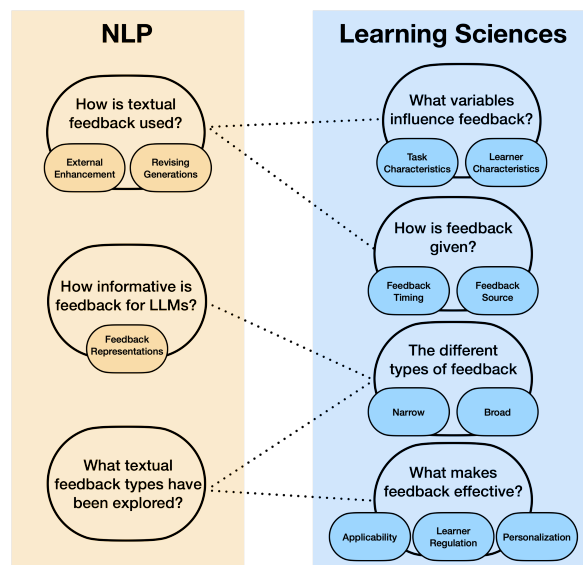


Figure 1: Connecting feedback research in NLP to foundations of feedback in the Learning Sciences.

Bai et al., 2022b; Saunders et al., 2022; Madaan et al., 2023) feedback, all variants have been shown to be successful in several metrics — from encouraging honest behaviors, to reducing toxicity, to being preferred by evaluators (Ouyang et al., 2024). Other approaches, such as imitation learning (Li et al., 2017; Kreutzer et al., 2018; Hancock et al., 2019; Scheurer et al., 2022), and feedback modeling (Weston, 2016; Li et al., 2017; Hancock et al., 2019; Xu et al., 2023; Liu et al., 2023) have had similar success. Feedback has thus emerged as an important source of information for model improvement and evaluation, directing them toward desired objectives and behaviors (Fernandes et al., 2023).

Despite the observed benefits of learning from feedback, a systematic study of what constitutes helpful feedback has so far remained absent. For example, RLF requires a Reward Model (RM) to be trained on numerical or ranking-based feedback data (Rafailov et al., 2023) — a format limited in the amount of information it conveys (Wu et al.,

2023). To counteract this limitation, works have begun exploring Natural Language Feedback (NLF; Weston, 2016; Madaan et al., 2023; Wu et al., 2023). However, these works rely on “intuitive guesses” about what constitutes useful feedback — leading different works to explore distinct conceptualizations of NLF, preventing a systematic comparison.

To establish a more concrete foundation for learning from feedback, we survey comprehensive studies of feedback from the field of learning sciences, which investigate feedback as an essential component of instruction and learning. We unify feedback approaches for NLP into a common framework, grounded on our surveyed pedagogical foundations (as summarized in Figure 1). To this end, we first present the most relevant feedback-related models from the learning sciences (§3). Taking inspiration from these pedagogical models, we create a novel framework, FELT, that expansively maps the various features of the LLM feedback space to pedagogical foundations (§4), and identify both dimensions studied by previous works as well as several important aspects of feedback that remain underexplored. We then focus on the least explored one, *feedback content*, and introduce a novel comprehensive taxonomy systematizing both the content and the delivery of natural language feedback (§4.2), shining a light on the underspecification of current approaches to feedback, and suggesting promising areas of future research.

Contributions We (i) present a survey of pedagogical feedback formulations and models; (ii) organize the variables that influence feedback (and its processing) into a schematic framework, specifically adapted for LLMs; (iii) propose a general and extensive taxonomy of feedback content; and (iv) propose areas of future research based on gaps between our taxonomy and the current landscape of LLM research on feedback.

2 Feedback in NLP

In this section, we survey current conceptualizations and applications of feedback in NLP, before exploring the feedback models and perspectives developed in the domain of learning science (§3).

2.1 Current State of Feedback in NLP

The value of feedback is derived from the implicit knowledge it represents about human values and expectations, that would otherwise be extremely difficult to specify (Christiano et al., 2017). Feed-

back can assume different forms: numerical ratings, rankings, preferences, demonstrations, and textual information (either using a template or free-form text — structured and unstructured feedback, respectively). RLF methodologies usually collect either a numeric rating or a ranking for classifying the *quality* of an initial answer (often encouraging properties such as *helpfulness* and *honesty* while mitigating *harmfulness*; Askell et al., 2021). RLF may also leverage demonstrations to finetune LLMs in a supervised fashion before the RLF stage to reduce the search space (Ouyang et al., 2024; Bai et al., 2022b). RLF methods address the intractable problem of designing an appropriate loss function for training models to exhibit behaviors with no closed-form solution (Askell et al., 2021). Recent works have also started to incorporate feedback in a prompt-based manner (Schick et al., 2023; Madaan et al., 2023; Paul et al., 2024; Chen et al., 2024; Lin et al., 2024; Zhao et al., 2024).

Informativeness of Feedback However, the extent to which feedback formulations transmit their goal states remains unclear. For example, Instruct-GPT (Ouyang et al., 2024) is first finetuned on demonstration data, and then performs RLHF with a reward model trained using comparison data (i.e., pairs of ranked generations). This feedback is limited in the information it transmits. Scoring a demonstration A as “better” than demonstration B provides little information on the quality of A nor B,¹ nor on how either demonstration can be improved. Taking these limitations into account, RMs are likely to suffer from some degree of misalignment (Pan et al., 2022; Gao et al., 2023; Song et al., 2023). Recent works have started to acknowledge the limited information in these feedback formulations, recognizing them as unsuited for capturing critical information, such as different types of errors (Golovneva et al., 2023; Wu et al., 2023).

2.2 Natural Language Feedback for LLMs

The most commonly used feedback formulations, scalar and ranking feedback, are thus limited in the information they convey, motivating new methods that leverage natural language for more expressive feedback formulations (Fernandes et al., 2023).

How should feedback be provided to LLMs?

Certain works augment the model through data augmentation (Shi et al., 2024), external corrective

¹We note such a format also obfuscates any bias and disagreement that occurred in reaching such a judgment.

feedback (Tandon et al., 2022; Madaan et al., 2022; Shinn et al., 2023) or natural language patches (Murty et al., 2022). Another line of work introduces a secondary model, that either refines an original LLM’s answer (Scheurer et al., 2022; Welleck et al., 2023; Tandon et al., 2022), critiques it (Saunders et al., 2022; Paul et al., 2024) iteratively self-improves (Schick et al., 2023; Chen et al., 2024; Madaan et al., 2023; Yuan et al., 2024) or otherwise constrains the initial response of a model (Stephan et al., 2024). Beyond increasing the complexity of feedback using natural language, several of these approaches also target intermediate generations with their feedback, not the final outcome produced by the model, thereby increasing the number of feedback opportunities through multiple iterations (Lightman et al., 2024).

What information should feedback content convey? Shi et al. (2024) distinguish textual feedback depending on whether the feedback is formally provided to the model’s answer, or, remains in the dialog setting, where the user mentions they disliked the reply they received. SELF-REFINE (Madaan et al., 2023) argues that the quality of the generated feedback is critical, though they only compare their “actionable and specific” LLM-generated feedback against “generic feedback” and no feedback at all. Wu et al. (2023) propose the introduction of finer-grained feedback — and of three different error types: factual incorrectness, irrelevance, and information incompleteness. Despite its impressive performance, the feedback exploration in this work is limited at only three specific types, and only uses preference rankings. Finally, Weston (2016) conducted the most thorough exploration of NLF to date, with 10 different dialogue-based supervision modes, which represent different interaction and feedback types. However, these often overlap information-wise, limiting its conclusions.

What is missing? Various works conceptualize feedback in completely different ways. No work has taken up a true mapping of the feedback space, identified the different types of information that can be encoded in NLF, and allowed for an exploration of different feedback components and their effectiveness. To address this, we look to the learning sciences, which studies feedback as an integral component of human learning. We survey their different conceptualizations of feedback (§3), and unify them under a new framework, FELT, adapted specifically for LLMs (§4).

3 Feedback in Learning Sciences

To construct a comprehensive, theory-grounded model of feedback that addresses the limited exploration of NLF in NLP, we build off the work of Lipnevich and Panadero (2021), who conducted a systematic review of 15 relevant and influential works on feedback models research in education. We provide a brief overview of the key points of each of these works and use them to subsequently propose a framework for the feedback ecosystem (§4) and a taxonomy for feedback content (§4.2).

3.1 What is feedback?

Many works agree that feedback is, or contains, information provided to a learner. While studies may disagree on the breadth or specificity required for feedback, and other limitations, a definition (adopted throughout this paper) emerges from their consensus:² *feedback is any task-relevant information given to a learner (content), by any possible feedback-generating agent (source).*

3.2 What constitutes effective feedback?

Kluger and DeNisi (1996) showed feedback was detrimental to a learner’s performance in 38% of analyzed cases. Three main requirements for helpful feedback emerge from previous work: *applicability*, *learner regulation*, and *personalization*. These requirements are directly related to feedback content, which we explore in §4.2 for LLMs.³

Applicability *Feedback should be actionable, allowing the learner to achieve a desired target performance.* Sadler (1989) suggests that feedback needs to identify a target performance, compare the learner’s current performance to it, and engage in actions to reduce that difference. Similarly, Hattie and Timperley (2007) indicate that effective feedback needs to answer three questions: where the learner is going (the goal), how they can get there, and where to go next. Other works extend these definitions of effective feedback by including elements such as motivational and metacognitive aspects (Nicol and Macfarlane-Dick, 2006) or aspects of teaching (e.g., lesson design; Evans, 2013).

Learner Regulation *Effective feedback produces a positive response in the learner.* Kluger and

²For an overview of all the different definitions of feedback discussed, please see Appendix C.

³Looking forth, *Applicability* corresponds to dimensions of *applicability of instructions*, *learner regulation to purpose*, and *personalization* is a product of the combination of all dimensions, allowing a fine-grained customization of feedback.

DeNisi (1996) argue that, in response to feedback, a learner’s attention will be directed to one of three levels: how to solve the task, the task as a whole, or meta-task processes (processes the learner performs while doing the task). Other works note that effective feedback also enhances self-regulated learning behaviors (Nicol and Macfarlane-Dick, 2006; Evans, 2013). Narciss and Huth (2004); Narciss (2008) extend this definition by arguing that feedback can have three distinct types of impact: influence on the learner’s cognitive abilities, their metacognitive skills, or their motivation and self-regulation. Lipnevich et al. (2016) defend that when a student receives feedback, they produce cognitive and affective responses, judging the task, their level of control, and the feedback. This produces a behavioral reaction, influencing their performance and learning. Similarly, Panadero and Lipnevich (2022) state that feedback impacts both the students’ performance and learning as well as their affective processes and self-regulation.

Personalization *Different types of feedback are best suited for different learner characteristics (Mason and Bruning, 2001).* The learner’s individual characteristics will also directly impact how they process feedback (Lipnevich et al., 2016).

3.3 What are the characteristics of feedback?

A large body of research has attempted to systematically categorize feedback based on its *content*, how it is given (*timing, source*), and the variables influencing it (*task, learner*). Works that systematically categorize feedback can be broadly divided into two groups: taxonomies focusing only on the content of feedback, and taxonomies taking into account the whole ecosystem.⁴

Content Works in this category focus on the characteristics of the content only. Kulhavy and Stock (1989), for example, model feedback through a *verification* component, which is a simple discrete classification of the answer as correct or incorrect, and an *elaboration* component, which contains all other information. Other works (Hattie and Timperley, 2007; Panadero and Lipnevich, 2022) suggest three categories for classifying feedback: (i) addressing the learner’s performance goal, (ii) addressing the learner’s current performance, and (iii) addressing the next steps the learner should undertake. Our feedback content taxonomy draws primarily from

⁴Appendix C presents a more thorough definition of proposed feedback categories in the learning sciences.

the analysis and adaptation of these works to the NLP domain. Additionally, several categorizations include feedback about mistakes, from which we derive the Error component of FELT (§4).

Ecosystem Other works suggest a more comprehensive categorization of feedback including the whole feedback ecosystem. For example, several works propose different feedback categories that consider characteristics of the learner (student proficiency, prior knowledge) and the task (difficulty) (Mason and Bruning, 2001; Narciss and Huth, 2004; Narciss, 2008). We further explore this categorization in §3.3.2.

3.3.1 How is feedback given?

Apart from its content and ecosystem, feedback is also characterized by how it is given. Two main components have emerged in the literature:

Source Feedback can be given by different sources (e.g., teachers, peers, or even the learner themselves). In their systematic review, Lipnevich and Panadero (2021) found that seven out of the 15 considered models view the source as an important characteristic of feedback. An additional three works distinguish feedback generated by an external source from self-feedback.

Timing Previous work has differentiated between immediate and delayed feedback. While early work (Bangert-Drowns et al., 1991) found delayed feedback to be more effective, more recent works (Mason and Bruning, 2001) argue that the optimal timing of feedback depends on learner characteristics. For example, Hattie and Timperley (2007) state that the most beneficial timing depends mainly on the complexity of the task. Complex tasks benefit from delayed feedback as they allow the learner to properly process the task. Other works (Narciss, 2008) focus mainly on the learner characteristics, arguing that as long as the learner possesses the metacognitive skills required to spot and address mistakes, feedback should be delayed.

We incorporate both feedback source and timing in our framework presented in §4.⁵

3.3.2 What variables influence feedback?

Feedback cannot only be categorized according to its content and the way it is given, but is also characterized by the ecosystem surrounding it. Many authors mention the challenge of determining the

⁵Other points of contention exist, such as feedback valence, but we consider valence as an element of the feedback’s content, and as such discuss it only in §4.2.

optimal feedback type in isolation. Instead, characteristics of the *task* and *learner* must be taken into account when giving feedback. We discuss their pedagogical definitions in this section, and subsequently adapt both *Task* and *Learner* specifically to LLMs in our framework, FELT (§4).

Task The task’s characteristics influence the optimal timing of feedback (§3.3.1) as well as its content. Mason and Bruning (2001) take into account the complexity of the task when choosing the most suitable feedback for a given setting and Narciss and Huth (2004); Narciss (2008) incorporate the task and the instructional content and goals into the instructional factors that affect feedback. The nature of the task (e.g., closed vs. open-answer) also influences feedback content and processing (Lipnevich and Panadero, 2021).

Learner Mason and Bruning (2001) consider student achievement and prior knowledge as important factors impacting feedback optimality. Nicol and Macfarlane-Dick (2006) expand prior knowledge and proficiency into *domain knowledge* and *strategy knowledge* (along with *motivational beliefs*), updated upon each attempt through both internal and external feedback. Narciss and Huth (2004); Narciss (2008) flesh out the learner characteristics further, including learning goals and motivation. Lipnevich et al. (2016) identify “personality, general cognitive ability, receptivity to feedback, prior knowledge, and motivation” as key learner characteristics that impact feedback processing.

Other works focus on the learner’s feedback processing mechanisms. Kluger and DeNisi (1996) propose three processing levels: details about how to solve the task, the task itself as a whole, and meta-task processes. In contrast, Hattie and Timperley (2007) present four levels of feedback processing: (1) *task level*, how well the tasks are understood and performed, (2) *process level*, the process needed to achieve the *task level* understanding and performance, (3) *self-regulation level*, the self-monitoring and the direction and regulation of the learner’s actions, and (4) *self level*, personal evaluations about the learner as a whole (which are usually uninformative about the task and can be very hard to change, e.g., the quality of being “a good student”). Current LLMs do not seem capable of more metacognitive levels of feedback processing.

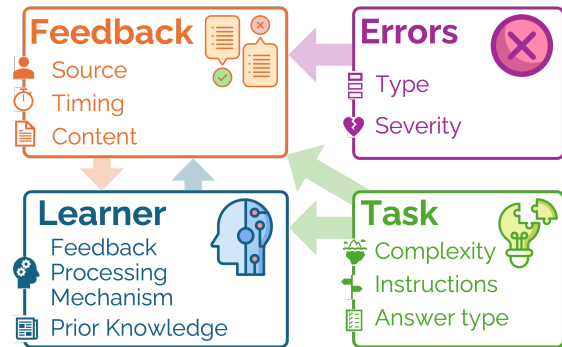


Figure 2: **FELT**, the feedback ecosystem adapted for LLMs. The feedback’s characteristics, the errors, the task, and the learner all influence both the feedback and its subsequent processing by the learner. Various interactions occur between the task, the learner, and the feedback, such as the task complexity and prior knowledge of the learner affecting the timing of the feedback. Appendix A presents a more comprehensive overview of the interactions present in the framework.

4 FELT: Unifying The Two Worlds

In the pedagogical research surveyed in the previous section (§3), feedback emerges as both a complex ecosystem, and a rich, but systematized, information source. In this section, we consolidate different pedagogical feedback models, and conceptualize a novel feedback framework adapted for LLMs, FELT (Feedback, Errors, Learner, Task), displayed in Figure 2. We briefly analyze each of its components, and dedicate a specific section (§4.2) to the formalization of feedback content.

4.1 The Big Picture

By consolidating the dimensions presented in Section 3, we developed FELT, which incorporates the content and delivery of feedback, as well as the rest of the ecosystem: the task, learner, and errors.

Task The task reflects the work assigned to the LLM. It has three attributes: *complexity* (i.e., difficulty level), *instructions*⁶ given, and the task’s *answer type* — which we reduce to being closed-answer (where there is a single correct answer or a finite set of them) or open-ended. When ported to LLMs, the task component can be reduced to data, as the training data will reflect the distribution of tasks that must be learned. Unsurprisingly, the impact of alignment data on model behavior is an active area of research in the NLP community (Fan

⁶Instructions can pertain to a description of the task, to specific criteria the model must satisfy, or to further directions, e.g., how much the model must adhere to any given feedback.

et al., 2019; Ethayarajh et al., 2022; Wang et al., 2023; Taori et al., 2023; Guo et al., 2023; Lambert et al., 2023; Köpf et al., 2023; Bai et al., 2022a).

Learner The learner in FELT represents the LLM itself, represented by two components: *prior knowledge* (which is dependent on the task), and the *feedback processing mechanism*.⁷ *Prior knowledge* is captured by the model’s size, pretraining data and pretraining method. Model size is suggested to be linked to the model’s ability to effectively leverage feedback (Scheurer et al., 2022; Bai et al., 2022b). The pretraining data, as well as how the model was trained, similarly encode its initial ability to tackle a given task. The *feedback processing mechanism* reflects the method used to transmit feedback to the model, naturally influencing how the feedback is subsequently processed. For example, training objectives necessarily influence how models process and incorporate feedback. Fernandes et al. (2023) identify three common feedback integration mechanisms: feedback-based imitation learning, joint-feedback modeling, and reinforcement learning (which can be generally extended to include other recent training methods, such as DPO, Rafailov et al., 2023). In addition, we also consider feedback provided using in-context learning (ICL; Brown et al., 2020; Madaan et al., 2022; Zhao et al., 2024). Much like the *Task*, the *Learner* (i.e., the models and algorithms used for learning from feedback) is an active area of NLP research.

Errors Feedback may relay information on where the learner is failing, which requires understanding the possible error modes for a given task, and which ones the learner is likely in. For example, guessing and committing systematic reasoning mistakes reflect differing degrees of understanding. Moreover, not all errors should be treated the same. Wu et al. (2023) find that assigning distinct reward models to specific error modes (e.g., conciseness, factuality, relevance) improved performance over using a single RM, clarifying the importance of modeling the error space – another active area of research (Golovneva et al., 2023; Paul et al., 2024;

⁷Several pedagogical dimensions for human feedback processing, discussed in §3, are not currently applicable to LLMs — such as the *learner’s regulation*, the learner’s *affective processes* and multiple *motivations* (e.g., both succeeding at the task and being seen as competent, which can be a deterrent to success if the learner believes they might fail), the learner’s perception of their *control over the task*, autonomous *receptivity to feedback*, and higher levels of feedback processing, including *meta-task* and *self level* learning — and as such do not have a direct representation in FELT’s current form.

Murugesan et al., 2023; Mishra et al., 2024).

Feedback Pedagogy distinguishes three important variables for feedback: content, source, and timing. The *content* captures both the type and the format of the information provided in a feedback message, two aspects we will explore fully in Section 4.2. The *source* of feedback typically identifies whether the feedback stems from an authority figure or a peer. In the context of LLMs, the source could be either a human or a separate AI model and receiving this information might lead LLMs to behave differently. However, current prevailing LLM architectures can only perceive the source of information by having it conveyed as text input to the model, reducing it to a dimension of content.

Timing specifies whether feedback is provided immediately or after a delay. Certain types of prompting can be formulated as variations in timing. For example, Metacognitive reasoning (Wang and Zhao, 2024) prompts the model to answer a question, critically evaluate its answer, possibly revise it, and assess its confidence in it. In contrast, Chain-of-Thought (CoT; Wei et al., 2024) does not elicit any post-answer reasoning. Coupled with feedback learning methods such as RLF or DPO, this “padding” between the answer and the reward signal might affect which prompting strategy leads to the best performance. Delayed feedback also has the potential for deriving multiple model updates from a single initial data point. From an answer and a reflection generated by an LLM (i.e., an attempt to identify and correct errors on that answer), and a piece of feedback, a model can learn not only whether the answer is correct, but also whether its reflection is adequate — applicable both in ICL and traditional training (e.g., RLF, DPO) settings.

4.2 Zooming in on Feedback Content

Of the four dimensions in FELT, feedback emerges as the least systematized, particularly when it is given through text or natural language. As briefly mention in Section 3.3 and detailed in Appendix C, several works have sought out to characterize feedback focusing only on its content. We draw from this research to define a novel taxonomy for feedback content, for which all dimensions can be controlled for and adjusted by the feedback provider.

Defining feedback content We distill four non-overlapping areas from categorizations of feedback in pedagogy research (§3):

1. **Learner status.** Situates the learner’s per-

- formance. May indicate what the learner got right, what mistakes the learner made, or both.
2. **Goal.** Demonstrates the target performance for the task by providing either the correct answer or an example solution.
 3. **Procedural.** Provides instructions for the learner to follow on a subsequent attempt. These recommendations can be specific to the task, be general problem-solving and metacognitive instructions, or both.
 4. **Peripheral.** Supplementary information not directly related to the above three areas. Peripheral feedback can include the clarification of the task (without any instructions), or the elaboration of concepts relevant to the task.

Modulating feedback content Once the decision of *what* content to include or omit in the feedback has been taken, *how* the feedback will be given must also be decided. As such, by adapting pedagogical models as well as incorporating entirely new axes exclusive to LLMs, we propose 10 dimensions that capture the form feedback takes, and that can be controlled by the feedback giver:

1. **Granularity**: measure of detail with which the feedback addresses the original answer.⁸
2. **Applicability of instructions**: outlines if the feedback contains instructions, and how applicable those instructions are for the learner and their current approach to solving the task.
3. **Answer coverage**: registers how much of the learner’s answer is reflected in the given feedback. The feedback could be independent of the answer, relate only to parts of the answer (e.g., focusing on a particular mistake), or take the complete answer into consideration.
4. **Target coverage**: indicates how much of the target performance is being considered when generating the feedback. *Goal* and *Procedural* types of feedback will both have at least some degree of *target coverage*.
5. **Criteria**: denotes which criteria the answer is being evaluated on: global evaluation, specific dimensions (e.g., fluency, engagement, etc.), or, alternatively, no dimensions (the answer is not being evaluated).

⁸For an open-answer example task, feedback might range from global meta-feedback, to task-specific, to paragraph-level, to sentence-level, to word-level, to token-level feedback. Appendix B provides several examples.

6. **Information novelty**: indicates the degree to which learner already had access to the information in the feedback, ranging from all information being previously known, to some information being unknown to the learner, to all information being novel for the learner.
7. **Purpose**: measures whether the feedback is being given to improve the learner’s performance or to clarify the task.⁹
8. **Style**: captures the type of language used to transmit the feedback to the learner, which can range from simple, direct sentences to verbose and terminology-heavy language.
9. **Valence**: indicates whether the feedback is positive (signaling achievement) or negative (signaling need for improvement).
10. **Mode**: captures whether the feedback is uni- or multimodal.¹⁰

The ability to modulate these ten dimensions enables practitioners to specifically craft feedback for their use cases and learning frameworks. Moreover, our feedback taxonomy is general and transferable across domains. While tasks can vary significantly in both type (e.g., coding, mathematics, story telling, common sense) and necessary skills (e.g., reasoning, planning, retrieval), the same type of feedback applies to all of them.

5 Realizing the Promise of Feedback

FELT provides a springboard for exploring the entire feedback ecosystem. To understand what makes feedback effective for LLMs, we briefly propose some ideas for future investigation (and categorize ideas from previous studies) that demonstrate the important role that each of the different components of FELT play in this exploration. Figure 3 presents a visual summary of the proposals presented in this section.

Task. Feedback might influence model performance differently depending on the **goal formulation** for the LLM (e.g., asking the model to extract entities from text v.s. asking the model to identify

⁹In a pedagogical setting with human learners, other purposes are possible, such as regulating the student’s emotions and motivation, but we do not consider these for LLMs.

¹⁰Multimodal feedback is naturally more suited for multimodal tasks. For example, in an instance segmentation task, the correct (visual) answer could be provided alongside textual feedback on mistakes and how to correct them.

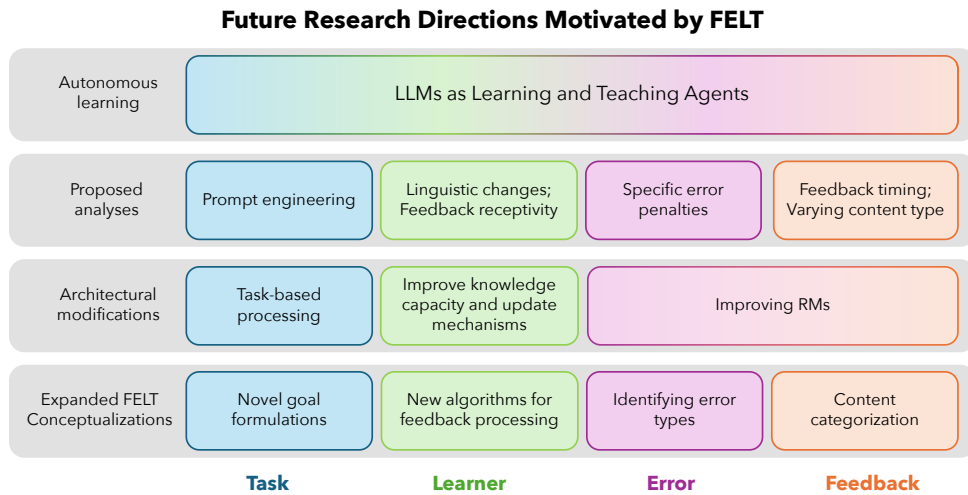


Figure 3: **Visual summary of several future research directions motivated by FELT**, for each of its dimensions and across several research axes. The textual counterparts of each bubble are highlighted in their respective color (blue, green, purple, orange) in the main body of Section 5.

their start and end positions). **Task-specific processing** (e.g., LoRA; Hu et al., 2022) would allow models to adapt to a wider range of tasks.

Learner. To explore the degree to which alignment extends beyond **changing the linguistic style** of AI models (Lin et al., 2024), LLMs can be modeled as different types of learners — whether by (1) controlling their **knowledge capacity**, i.e., creating learners both with and without prior knowledge through data interventions during pretraining or finetuning, or external augmentation; and (2) implementing **novel update mechanisms** that can capture richer feedback (e.g., TextGrad; Yuksekgonul et al., 2024). Uncertainty awareness has also been shown to improve model alignment (Wang et al., 2024). Exploring the **model’s receptiveness to feedback** conditioned on its uncertainty may enable better understanding of what leads models to reject their own parametric knowledge.

Errors. After **mapping the different error types** a model might make, it becomes possible to implement **different penalties for these error types**, which might promote different model behaviors (Wu et al., 2023). This flexibility will yield a better understanding of which degrees of penalization lead to a given behavior, as well as the degree to which this penalization should be aligned with the actual severity that the error merits.

Feedback. By delaying the **timing of the feedback** and first asking the model for a reflection, it is possible to multiply the datapoints the model is trained on (generating feedback not only on the correctness of the answer but also on that of the

model’s reflections). This approach has, to the best of our knowledge, not yet been pursued.

Categorizing Feedback A limitation of prior works on NLF is the indiscriminate treatment of different types of feedback (shown in Table 1). The four types of feedback content — *learner status*, *goal*, *procedural*, and *peripheral* — directly address this issue, allowing more systematic studies of feedback. The natural next step is to provide models with **various kinds of feedback** and compare the impact of the different types of information and of its delivery.

Prompt Engineering ICL is an active area of research in NLP, including in model alignment (Lin et al., 2024; Zhao et al., 2024). Several dimensions of our feedback content taxonomy have been shown to contribute to either aligning LLMs to specific desired behaviors, or to increasing its performance in a given task — suggesting both the utility of these content dimensions, and hinting at many as of yet unexplored approaches that vary them in novel ways. Prompts with a high degree of **applicability of instructions** and with a focus on *procedural* types of information are common, be it asking for a rationale (Wei et al., 2024; Wang and Zhao, 2024; Yao et al., 2023), or more task-specific instructions (Wang and Zhao, 2024; Madaan et al., 2023). Prompting the model with *goal* information is also popular, most frequently done using few-shot prompting (Brown et al., 2020). **Valence** is often introduced either from (explicit or implicit) human feedback (Shi et al., 2024; Pang et al., 2024)

Content Type	Previous Works
Learner Status	Weston (2016), Tandon et al. (2022), Scheurer et al. (2022), Saunders et al. (2022), Shi et al. (2024), Welleck et al. (2023), Wu et al. (2023), Shinn et al. (2023), Chen et al. (2024), Paul et al. (2024), Madaan et al. (2023), Lightman et al. (2024), Yoon et al. (2024)
Goal	Weston (2016), Saunders et al. (2022), Shi et al. (2024)
Procedural	Weston (2016), Tandon et al. (2022), Saunders et al. (2022), Shi et al. (2024), Welleck et al. (2023), Madaan et al. (2023), Murty et al. (2022), Schick et al. (2023)
Peripheral	Tandon et al. (2022), Scheurer et al. (2022), Chen et al. (2024), Madaan et al. (2023)

Table 1: **Past works using textual or otherwise augmented feedback.** Even in the same work, feedback such as “write a step-by-step reasoning trace before the solution” is often treated as equivalent to “should be contextually relevant and easy to check.” As made evident by our taxonomy, however, they are markedly different, as the former provides *procedural* instructions and the latter *peripheral* information clarifying the task.

or by criticizing a preliminary answer (Bai et al., 2022b; Paul et al., 2024; Yao et al., 2023; Yuan et al., 2024), providing the model with information on its *learner status*. The *style*, or linguistic properties, of prompts can lead to significant performance variations (Leidinger et al., 2023). Finally, while *Information Novelty* is hard to measure for LLMs, recent research has investigated what causes LLMs to be faithful to the novel information rather than their encoded knowledge (Longpre et al., 2021), an important quality to understand and predict model behavior (Zhou et al., 2023; Yin et al., 2023). All these dimensions underlie a wide range of instruction formulations for LLMs, enabling practitioners to taxonomize existing prompting strategies, and uncover novel approaches based on different combinations.

Improving Reward Models Textual feedback can be integrated into the training loop of reward models, whether in a traditional training pipeline (i.e., RLF) or in inference-time learning (e.g., Inference-Time Policy Adapters; Lu et al., 2023). For example, textual feedback about a policy model’s output could be provided in addition to its output to the reward model, or, following Wu et al. (2023), the model can be rewarded according to a more targeted set of preferences. It is possible to go further still by varying the *granularity* of the feedback, allowing targeted rewards for different parts of the model response to be generated. A similar effect can be achieved by considering a set of different *criteria* on which to judge model performance. RMs themselves can be extended, by adding adapters for different criteria or using a Mixture-of-Experts (MoE) to better model these fine-grained preferences. The exploration of these dimensions allows for more expressive RMs, and

possibly modeling non-deterministic preferences.

LLMs as Learning and Teaching Agents By defining the various factors that impact learning from feedback, FELT allows the creation of a *environments* to explore and optimize feedback for LLMs. Multiple LLM agent learners can be simulated in this environment, each receiving differently formulated feedback to guide them toward their goals. This environment could enable the study of the learning process from initial exposure to data to more complex interaction with other learner agents. The mapping provided by FELT similarly enables the deployment of LLMs as teachers by identifying the necessary components needed for better feedback dataset construction, making feedback more adapted and personalized for learner models.

6 Conclusion

We survey the most influential feedback models from the learning sciences, creating a novel framework, FELT, as well as a taxonomy of natural language feedback (NLF). Both enable the systematic exploration of feedback, allowing for objective conclusions on the impact of a piece of feedback and the optimization of NLF. Beyond a survey of the space of natural language feedback, we explore how both FELT and the fine-grained feedback content dimensions underlie many existing techniques in natural language feedback in AI, making specific recommendations for enhancing feedback formulations for a wide range of proposed topics.

Limitations

FELT, our proposed framework to capture the full feedback ecosystem, is only theoretically grounded at the moment. Certain aspects of our taxonomy,

such as the impact of feedback timing, need to be empirically assessed for LLMs. Similarly, while we conjecture that all 10 dimensions of the feedback content taxonomy will impact how models react to feedback, this inference has not yet been empirically confirmed. However, our grounding of the LLM feedback space to pedagogical principles is meant to provide a broad framework for organizing feedback research, with each component available for empirical validation by the research community. Components of our framework that end up failing empirical validation perhaps indicate areas where LLMs differ from human learning.

Another limitation of our work is that both pedagogical and NLP conceptualizations and results we discussed in this paper were conducted in English settings. While we expect FELT and our feedback content taxonomy to generalize to other languages, the same feedback content might affect models differently depending on the language with which it is given. This dimension must be taken into account by future work. Finally, any study into how to make feedback more effective has the potential to contribute to the jailbreaking of LLMs or other purposefully malicious changes in its behavior. However, we also note that better understanding the components of feedback that make it effective will enable researchers to develop models that are better aligned to their original goals and perhaps more robust to these types of attacks.

Acknowledgements

We gratefully acknowledge the support of the Swiss National Science Foundation (No. 215390), Innosuisse (PFFS-21-29), the EPFL Science Seed Fund, the EPFL Center for Imaging, Sony Group Corporation, and the Allen Institute for AI.

References

- Amanda Askell, Yuntao Bai, Anna Chen, Dawn Drain, Deep Ganguli, Tom Henighan, Andy Jones, Nicholas Joseph, Ben Mann, Nova DasSarma, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Jackson Kernion, Kamal Ndousse, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, and Jared Kaplan. 2021. [A general language assistant as a laboratory for alignment](#). *arXiv preprint arXiv:2112.00861*.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. 2022a. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *arXiv preprint arXiv:2204.05862*.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. 2022b. [Constitutional ai: Harmlessness from ai feedback](#). *arXiv preprint arXiv:2212.08073*.
- Robert L. Bangert-Drowns, Chen-Lin C. Kulik, James A. Kulik, and MaryTeresa Morgan. 1991. [The Instructional Effect of Feedback in Test-like Events](#). In *Review of Educational Research*, volume 61, pages 213–238. Sage Publications, Inc.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Deborah L. Butler and Philip H. Winne. 1995. [Feedback and Self-Regulated Learning: A Theoretical Synthesis](#). In *Review of Educational Research*, volume 65, pages 245–281. Sage Publications, Inc.
- David Carless and David Boud. 2018. [The development of student feedback literacy: enabling uptake of feedback](#). In *Assessment & Evaluation in Higher Education*, volume 43, pages 1315–1325. Routledge.
- Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. 2024. [Teaching large language models to self-debug](#). In *The Twelfth International Conference on Learning Representations*.

- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martić, Shane Legg, and Dario Amodei. 2017. [Deep reinforcement learning from human preferences](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Kawin Ethayarajh, Yejin Choi, and Swabha Swayamdipta. 2022. Understanding dataset difficulty with \mathcal{V} -usable information. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 5988–6008. PMLR.
- Carol Evans. 2013. [Making sense of assessment feedback in higher education](#). In *Review of Educational Research*, volume 83, pages 70–120. Sage Publications.
- Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. 2019. [ELI5: long form question answering](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3558–3567. Association for Computational Linguistics.
- Patrick Fernandes, Aman Madaan, Emmy Liu, António Farinhas, Pedro Henrique Martins, Amanda Bertsch, José G. C. de Souza, Shuyan Zhou, Tongshuang Wu, Graham Neubig, and André F. T. Martins. 2023. [Bridging the gap: A survey on integrating \(human\) feedback for natural language generation](#). In *Transactions of the Association for Computational Linguistics*, volume 11, pages 1643–1668. MIT Press.
- Adam Fisch, Jacob Eisenstein, Vicky Zayats, Alekh Agarwal, Ahmad Beirami, Chirag Nagpal, Pete Shaw, and Jonathan Berant. 2024. [Robust preference optimization through reward model distillation](#). *arXiv preprint arXiv:2405.19316*.
- Leo Gao, John Schulman, and Jacob Hilton. 2023. [Scaling laws for reward model overoptimization](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 10835–10866. PMLR.
- Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland, Michal Valko, and Daniele Calandriello. 2024. [A general theoretical paradigm to understand learning from human preferences](#). In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pages 4447–4455. PMLR.
- Olga Golovneva, Moya Peng Chen, Spencer Poff, Martín Corredor, Luke Zettlemoyer, Maryam Fazel-Zarandi, and Asli Celikyilmaz. 2023. [ROSCOE: A suite of metrics for scoring step-by-step reasoning](#). In *The Eleventh International Conference on Learning Representations*.
- Biyang Guo, Xin Zhang, Ziyuan Wang, Minqi Jiang, Jinran Nie, Yuxuan Ding, Jianwei Yue, and Yupeng Wu. 2023. How close is chatgpt to human experts? comparison corpus, evaluation, and detection. *arXiv preprint arxiv:2301.07597*.
- Braden Hancock, Antoine Bordes, Pierre-Emmanuel Mazare, and Jason Weston. 2019. [Learning from dialogue after deployment: Feed yourself, chatbot!](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3667–3684, Florence, Italy. Association for Computational Linguistics.
- John Hattie and Helen Timperley. 2007. [The Power of Feedback](#). In *Review of Educational Research*, volume 77, pages 81–112.
- Jiwoo Hong, Noah Lee, and James Thorne. 2024. [ORPO: Monolithic preference optimization without reference model](#). *arXiv preprint arXiv:2403.07691*.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Avraham N. Kluger and Angelo DeNisi. 1996. [The effects of feedback interventions on performance: A historical review, a meta-analysis, and a preliminary feedback intervention theory](#). In *Psychological Bulletin*, volume 119, pages 254–284.
- Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi Rui Tam, Keith Stevens, Abdullah Barhoum, Duc Minh Nguyen, Oliver Stanley, Richárd Nagyfi, Shahul ES, Sameer Suri, David Alexandrovich Glushkov, Arnav Varma Dantururi, Andrew Maguire, Christoph Schuhmann, Huu Nguyen, and Alexander Julian Mattick. 2023. [OpenAssistant Conversations - democratizing large language model alignment](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Tomasz Korbak, Kejian Shi, Angelica Chen, Rasika Bhalerao, Christopher L. Buckley, Jason Phang, Samuel R. Bowman, and Ethan Perez. 2023. Pre-training language models with human preferences. In *Proceedings of the 40th International Conference on Machine Learning*. JMLR.
- Julia Kreutzer, Shahram Khadivi, Evgeny Matusov, and Stefan Riezler. 2018. [Can neural machine translation be improved with user feedback?](#) In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers)*, pages 92–105, New Orleans - Louisiana. Association for Computational Linguistics.
- Raymond W. Kulhavy. 1977. [Feedback in Written Instruction](#). In *Review of Educational Research*, volume 47, pages 211–232. Sage Publications, Inc.

- Raymond W. Kulhavy and William A. Stock. 1989. [Feedback in written instruction: The place of response certitude](#). In *Educational Psychology Review*, volume 1, pages 279–308.
- Nathan Lambert, Lewis Tunstall, Nazneen Rajani, and Tristan Thrush. 2023. [Huggingface h4 stack exchange preference dataset](#).
- Alina Leidinger, Robert van Rooij, and Ekaterina Shutova. 2023. [The language of prompting: What linguistic properties make a prompt successful?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9210–9232, Singapore. Association for Computational Linguistics.
- Jiwei Li, Alexander H. Miller, Sumit Chopra, Marc’Aurelio Ranzato, and Jason Weston. 2017. [Dialogue learning with human-in-the-loop](#). In *International Conference on Learning Representations*.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2024. [Let’s verify step by step](#). In *The Twelfth International Conference on Learning Representations*.
- Bill Yuchen Lin, Abhilasha Ravichander, Ximing Lu, Nouha Dziri, Melanie Sclar, Khyathi Chandu, Chandra Bhagavatula, and Yejin Choi. 2024. [The unlocking spell on base LLMs: Rethinking alignment via in-context learning](#). In *The Twelfth International Conference on Learning Representations*.
- Anastasiya A. Lipnevich, David A. G. Berg, and Jeffrey K. Smith. 2016. [Toward a Model of Student Response to Feedback](#). In *Handbook of Human and Social Conditions in Assessment*. Routledge.
- Anastasiya A. Lipnevich and Ernesto Panadero. 2021. [A Review of Feedback Models and Theories: Descriptions, Definitions, and Conclusions](#). In *Frontiers in Education*, volume 6.
- Hao Liu, Carmelo Sferrazza, and Pieter Abbeel. 2023. [Chain of hindsight aligns language models with feedback](#). volume abs/2302.02676.
- Shayne Longpre, Kartik Perisetla, Anthony Chen, Nikhil Ramesh, Chris DuBois, and Sameer Singh. 2021. [Entity-based knowledge conflicts in question answering](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7052–7063, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ximing Lu, Faeze Brahman, Peter West, Jaehun Jung, Khyathi Chandu, Abhilasha Ravichander, Prithviraj Ammanabrolu, Liwei Jiang, Sahana Ramnath, Nouha Dziri, Jillian Fisher, Bill Lin, Skyler Hallinan, Lianhui Qin, Xiang Ren, Sean Welleck, and Yejin Choi. 2023. [Inference-time policy adapters \(IPA\): Tailoring extreme-scale LMs without fine-tuning](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6863–6883, Singapore. Association for Computational Linguistics.
- Aman Madaan, Niket Tandon, Peter Clark, and Yiming Yang. 2022. [Memory-assisted prompt editing to improve GPT-3 after deployment](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2833–2861, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. 2023. [Self-refine: Iterative refinement with self-feedback](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Beryl Mason and Roger Bruning. 2001. [Providing Feedback in Computer-based Instruction: What the Research Tells Us](#). In *Center for Instructional Innovation*, volume 15.
- Santosh Mathan and K. Koedinger. 2005. [Fostering the intelligent novice: Learning from errors with metacognitive tutoring](#). In *Educational Psychologist*, volume 40, pages 257 – 265.
- Yu Meng, Mengzhou Xia, and Danqi Chen. 2024. [SimPO: Simple preference optimization with a reference-free reward](#). *arXiv preprint arXiv:2405.14734*.
- Jacob Menick, Maja Trebacz, Vladimir Mikulik, John Aslanides, Francis Song, Martin Chadwick, Mia Glaese, Susannah Young, Lucy Campbell-Gillingham, Geoffrey Irving, and Nat McAleese. 2022. [Teaching language models to support answers with verified quotes](#). *arXiv preprint arXiv:2203.11147*.
- Abhika Mishra, Akari Asai, Vidhisha Balachandran, Yizhong Wang, Graham Neubig, Yulia Tsvetkov, and Hannaneh Hajishirzi. 2024. [Fine-grained hallucination detection and editing for language models](#). In *First Conference on Language Modeling*.
- Shikhar Murty, Christopher Manning, Scott Lundberg, and Marco Tulio Ribeiro. 2022. [Fixing model bugs with natural language patches](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11600–11613, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Keerthiram Murugesan, Sarathkrishna Swaminathan, Soham Dan, Subhajit Chaudhury, Chulaka Gunasekara, Maxwell Crouse, Diwakar Mahajan, Ibrahim Abdelaziz, Achille Fokoue, Pavan Kapanipathi, Salim Roukos, and Alexander Gray. 2023.

- MISMATCH: Fine-grained evaluation of machine-generated text with mismatch error types.** In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4485–4503, Toronto, Canada. Association for Computational Linguistics.
- Susanne Narciss. 2008. Feedback Strategies for Interactive Learning Tasks. In *Handbook of research on educational communications and technology*, pages 125–144. Routledge.
- Susanne Narciss and Katja Huth. 2004. **How to design informative tutoring feedback for multi-media learning.** volume Instructional design for multimedia learning, pages 181–195. Waxmann.
- David J. Nicol and Debra Macfarlane-Dick. 2006. **Formative assessment and self-regulated learning: a model and seven principles of good feedback practice.** In *Studies in Higher Education*, volume 31, pages 199–218. Routledge.
- OpenAI. 2023. **GPT-4 technical report.** *arXiv preprint arXiv:2303.08774*.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2024. Training language models to follow instructions with human feedback. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA. Curran Associates Inc.
- Alexander Pan, Kush Bhatia, and Jacob Steinhardt. 2022. **The effects of reward misspecification: Mapping and mitigating misaligned models.** In *International Conference on Learning Representations*.
- Ernesto Panadero and Anastasiya A. Lipnevich. 2022. **A review of feedback models and typologies: Towards an integrative model of feedback elements.** In *Educational Research Review*, volume 35.
- Richard Yuanzhe Pang, Stephen Roller, Kyunghyun Cho, He He, and Jason Weston. 2024. **Leveraging implicit feedback from deployment data in dialogue.** In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 60–75, St. Julian’s, Malta. Association for Computational Linguistics.
- Debjit Paul, Mete Ismayilzada, Maxime Peyrard, Beatriz Borges, Antoine Bosselut, Robert West, and Boi Faltings. 2024. **REFINER: Reasoning feedback on intermediate representations.** In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1100–1126, St. Julian’s, Malta. Association for Computational Linguistics.
- Jakub Prochazka, Martin Ovcari, and Michal Durinik. 2020. **Sandwich feedback: The empirical evidence of its effectiveness.** In *Learning and Motivation*, volume 71.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. **Direct preference optimization: Your language model is secretly a reward model.** In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Rajkumar Ramamurthy, Prithviraj Ammanabrolu, Kianté Brantley, Jack Hessel, Rafet Sifa, Christian Bauckhage, Hannaneh Hajishirzi, and Yejin Choi. 2023. **Is reinforcement learning (not) for natural language processing: Benchmarks, baselines, and building blocks for natural language policy optimization.** In *The Eleventh International Conference on Learning Representations*.
- Arkalgud Ramaprasad. 1983. **On the Definition of Feedback.** In *Behavioral Science*, volume 28, pages 4–13.
- D. Royce Sadler. 1989. **Formative assessment and the design of instructional systems.** In *Instructional Science*, volume 18, pages 119–144.
- Amir Saeidi, Shivanshu Verma, Aswin RRV, and Chitta Baral. 2024. **Triple preference optimization: Achieving better alignment with less data in a single step optimization.** *arXiv preprint arXiv:2405.16681*.
- William Saunders, Catherine Yeh, Jeff Wu, Steven Bills, Long Ouyang, Jonathan Ward, and Jan Leike. 2022. **Self-critiquing models for assisting human evaluators.** *arXiv preprint arXiv:2206.05802*.
- Jérémy Scheurer, Jon Ander Campos, Jun Shern Chan, Angelica Chen, Kyunghyun Cho, and Ethan Perez. 2022. **Training language models with language feedback.** *arXiv preprint arXiv:2204.14146*.
- Timo Schick, Jane A. Yu, Zhengbao Jiang, Fabio Petroni, Patrick Lewis, Gautier Izacard, Qingfei You, Christoforos Nalmpantis, Edouard Grave, and Sebastian Riedel. 2023. **PEER: A collaborative language model.** In *The Eleventh International Conference on Learning Representations*.
- Weiyan Shi, Emily Dinan, Kurt Shuster, Jason Weston, and Jing Xu. 2024. **When life gives you lemons, make cherryade: Converting feedback from bad responses into good labels.** In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3066–3082, Mexico City, Mexico. Association for Computational Linguistics.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik R Narasimhan, and Shunyu Yao. 2023. **Reflection: language agents with verbal reinforcement learning.** In *Thirty-seventh Conference on Neural Information Processing Systems*.

- Ziang Song, Tianle Cai, Jason D. Lee, and Weijie J. Su. 2023. [Reward collapse in aligning large language models: A prompt-aware approach to preference rankings](#). In *ICML 2023 Workshop The Many Facets of Preference-Based Learning*.
- Moritz Stephan, Alexander Khazatsky, Eric Mitchell, Annie S. Chen, Sheryl Hsu, Archit Sharma, and Chelsea Finn. 2024. [RLVF: Learning from verbal feedback without overgeneralization](#). *arXiv preprint arXiv:2402.10893*.
- Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. 2020. Learning to summarize from human feedback. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*, Red Hook, NY, USA. Curran Associates Inc.
- Niket Tandon, Aman Madaan, Peter Clark, and Yiming Yang. 2022. [Learning to repair: Repairing model output errors after deployment using a dynamic memory of feedback](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 339–352, Seattle, United States. Association for Computational Linguistics.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca.
- Yikun Wang, Rui Zheng, Liang Ding, Qi Zhang, Dahua Lin, and Dacheng Tao. 2024. [Uncertainty aware learning for language model alignment](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11087–11099, Bangkok, Thailand. Association for Computational Linguistics.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [Self-instruct: Aligning language models with self-generated instructions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13484–13508, Toronto, Canada. Association for Computational Linguistics.
- Yuqing Wang and Yun Zhao. 2024. [Metacognitive prompting improves understanding in large language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 1914–1926, Mexico City, Mexico. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2024. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA. Curran Associates Inc.
- Sean Welleck, Ximing Lu, Peter West, Faeze Brahman, Tianxiao Shen, Daniel Khashabi, and Yejin Choi. 2023. [Generating sequences by learning to self-correct](#). In *The Eleventh International Conference on Learning Representations*.
- Jason Weston. 2016. Dialog-based language learning. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS' 16*, page 829–837, Red Hook, NY, USA. Curran Associates Inc.
- Jeff Wu, Long Ouyang, Daniel M. Ziegler, Nisan Stiennon, Ryan Lowe, Jan Leike, and Paul Christiano. 2021. [Recursively summarizing books with human feedback](#). *arXiv preprint arXiv:2109.10862*.
- Zequ Wu, Yushi Hu, Weijia Shi, Nouha Dziri, Alane Suhr, Prithviraj Ammanabrolu, Noah A. Smith, Mari Ostendorf, and Hannaneh Hajishirzi. 2023. [Fine-grained human feedback gives better rewards for language model training](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Jing Xu, Megan Ung, Mojtaba Komeili, Kushal Arora, Y-Lan Boureau, and Jason Weston. 2023. [Learning new skills after deployment: Improving open-domain internet-driven dialogue with human feedback](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13557–13572, Toronto, Canada. Association for Computational Linguistics.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik R. Narasimhan. 2023. [Tree of thoughts: Deliberate problem solving with large language models](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Xunjian Yin, Baizhou Huang, and Xiaojun Wan. 2023. [ALCUNA: Large language models meet new knowledge](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1397–1414, Singapore. Association for Computational Linguistics.
- Eunseop Yoon, Hee Suk Yoon, Soohwan Eom, Gunsoo Han, Daniel Nam, Daejin Jo, Kyoung-Woon On, Mark Hasegawa-Johnson, Sungwoong Kim, and Chang Yoo. 2024. [TLCR: Token-level continuous reward for fine-grained reinforcement learning from human feedback](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 14969–14981, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Xian Li, Sainbayar Sukhbaatar, Jing Xu, and Jason E. Weston. 2024. [Self-rewarding language models](#). In *Forty-first International Conference on Machine Learning*.

Mert Yuksekgonul, Federico Bianchi, Joseph Boen, Sheng Liu, Zhi Huang, Carlos Guestrin, and James Zou. 2024. [Textgrad: Automatic "differentiation" via text.](#) *arXiv preprint arXiv: 2406.07496*.

Hao Zhao, Maksym Andriushchenko, Francesco Croce, and Nicolas Flammarion. 2024. [Is in-context learning sufficient for instruction following in llms?](#) *arXiv preprint arXiv:2405.19874*.

Wenxuan Zhou, Sheng Zhang, Hoifung Poon, and Muhao Chen. 2023. [Context-faithful prompting for large language models.](#) In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14544–14556, Singapore. Association for Computational Linguistics.

Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2020. [Fine-tuning language models from human preferences.](#) *arXiv preprint arXiv:1909.08593*.

A Components of the FELT Framework

The FELT framework introduced in Section 4 presents an important overview of all the factors that influence feedback and are in turn influenced by it. Figure 2 showcased a schematic overview of the FELT framework, integrating four distinct components: Feedback, Errors, Learner, and Task. In this appendix, we will outline more precisely each of the components of the FELT Framework, as well as the interactions between them.

A.1 Task

Typically, the task will be the first element to be defined.

Answer Type Understanding the answer is fairly straightforward – a task has a closed-answer if there is a finite set of correct answers, and an open-answer otherwise. Notably, tasks can contain both elements. For example the task "*Write a quality 4-paragraph short-story*" has both open- and closed-answer elements. There is no finite set of answer of what a quality story is, but whether a story has 4 paragraphs, or not, is a binary closed-answer task.

Complexity The difficulty level of a task is harder to define as some measure of relativity is involved. We suggest anchoring this measurement to the average adult human capabilities. A simple arithmetic task will thus be considered very easy, whereas researching and writing a doctoral thesis would be seen as hard.

Prompt Instructions The task instructions will be presented to the model at two distinct points in time: when first assigning the model this task, and when later providing feedback. With regards to the former, this element captures the degree to which the task is explained – is the model explicitly aware of all criteria it should satisfy? With regards to the second pass, when feedback is provided, this dimension pertains instead with the degree of freedom it gives the LLM – is the model forced to take the feedback into account, or can it consider only part of it, or even disregard it altogether if it deems it useless?

A.2 Learner

Either at the same time the task is defined or immediately after, the model to be tested will be chosen. The model choice influences two important features.

Prior Knowledge The prior knowledge captures the LLM’s abilities as a direct result of its size, training data, and training method. These, in turn, also reflect the model’s purpose (e.g., was it designed to be helpful, harmless, entertaining, etc.). The prior knowledge thus captures the model’s representation of the learner, and in its architecture and parametric knowledge, it encodes the LLM’s current abilities – or its proficiency – both in general and with regards to the specific task.

Feedback Processing Mechanism Mainly defined by the experimental setup, the mechanism by which the model process feedback can vary significantly, and not all of them are able to leverage the same level of information. Imitation learning, for example, can only leverage information which was positively evaluated. As stated in Section 4, we identify 4 main processing mechanisms, 3 of which alter the model’s parametric state – feedback-based imitation learning, joint-feedback modeling, and reinforcement learning, as defined in [Fernandes et al. \(2023\)](#) – and a fourth, non-parametric mode: in-context learning ([Brown et al., 2020](#)).

A.3 Errors

After both the task and learner are in place, the first pass of the experiment can be run, where the model will have its first attempt at solving the task. In this attempt, it is expected that the model will make some degree of mistakes – which have two important characteristics.

Error Type There are several possible types of errors, and their differences are significant. For example, an error made due to a guess only needs to provide the learner with the right information for it to be corrected, whereas a systematic error (for example, the mixing of British and American English spellings) will require a different, much more insistent, intervention. ROSCOE ([Golovneva et al., 2023](#)) proposes a taxonomy of step-by-step reasoning errors. While task dependent (i.e., there are grammar errors and arithmetic errors, rather than fully task independent failure modes), this taxonomy provides a good starting ground for the exploration of error types in NLP.

Error Severity Besides the type of error, it is also important to take the severity of the error into account. Stating that Marie Curie was a German philosopher and stating that she won one Nobel Prize in her lifetime are both factually inaccurate –

but one is a severe, complete hallucination, while the other omitted she actually won the Nobel Prize twice. The more severe the error, the stronger, more insistent, and more corrective the feedback should be.

A.4 Feedback

Finally, after the model has finished its first attempt at the task, producing some number of errors, feedback can be provided on this attempt.

Timing One easy to neglect aspect of feedback that pedagogy has shown to be impactful is timing – whether the feedback is provided immediately after a task is attempted or whether there is a delay between the two actions. There are differing opinions amongst education researchers, but how to make feedback content more effective through timing merit research in LLMs. For example, in line with [Mathan and Koedinger \(2005\)](#) and [Narciss \(2008\)](#)'s take on timing – delay feedback if the learner possesses metacognitive abilities that allow them to identify and possibly correct mistakes – we posit feedback will be more effective if, content-wise, it is preceded by information on the answer's correctness and mistakes' existence and only after this metacognitive priming is the rest of the information presented.

Content Section 4.2 explores feedback content in depth, presenting 10 impactful axes on which it can vary: length, granularity, applicability of instructions, answer coverage, criteria, information novelty, purpose, style, valence, and mode. It also presents a set of 9 emergent categories which, based on pedagogical research, we estimate to be the most promising one with regards to impact on revised model generations, and thus most deserving of further study.

Source Finally, it is also important to consider the source of feedback, which might be an authority, such as an expert, an average human, another LLM, a rule-based system, among others. Different sources will reflect different authority and reliability levels.

A.5 Interactions

With a clear understanding of all the components and sub-components of the FELT framework, we can explore the influences that exist between them.

Both the task complexity and the learner's prior knowledge can impact the ideal feedback timing –

be it delayed when the learner has metacognitive skills ([Narciss, 2008](#)) or enough task proficiency ([Mason and Bruning, 2001](#)) they can identify where the mistake occurred, or, for example, immediate if they don't ([Narciss, 2008](#)) or the task difficulty is low ([Mason and Bruning, 2001](#)).

With regards to the feedback content, the type of task ([Butler and Winne, 1995](#); [Kluger and DeNisi, 1996](#); [Mason and Bruning, 2001](#); [Lipnevich et al., 2016](#)) and both the error type and severity will have an impact ([Narciss and Huth, 2004](#); [Narciss, 2008](#)). The nature of the task (open or closed answer) will directly condition the feedback that can be given in response to the model's answer, as well as how difficult it will be to produce it. For example, generating the correct answer for a multiple choice quiz or a story writing task will be two very different endeavors. Similarly, it is impossible to provide response elaboration feedback on a single multiple choice question. The error type and severity will also influence the feedback content, as apart from directly dictating what mistakes verification and elaboration feedback can be given, they will also condition the ideal amount of detail and explanations to address the mistake at the most efficient level.

Finally, all aspects of feedback will influence the learner's feedback processing mechanism ([Kulhavy and Stock, 1989](#); [Sadler, 1989](#); [Bangert-Drowns et al., 1991](#); [Butler and Winne, 1995](#); [Kluger and DeNisi, 1996](#); [Narciss and Huth, 2004](#); [Nicol and Macfarlane-Dick, 2006](#); [Narciss, 2008](#); [Lipnevich et al., 2016](#); [Carless and Boud, 2018](#)). All three dimensions of feedback have evident potential to directly influence how the model processes them. The instruction's permissiveness to consider or discard feedback will also impact the learner's feedback processing mechanism. This processing is, of course, dependent on the specific processing mechanism employed, and while some might be indifferent to some of these components – like imitation learning, for example, which focuses exclusively on the feedback content – others will be sensitive to all, including the task's prompt instructions – such as in-context learning.

B Feedback content examples

To help concretize the 10 dimensions alongside which feedback content can be modulated, in this appendix we provide a few examples for each of these dimensions, with the exception of *mode*, which simply presents information in formats beyond text (e.g., images, audio, video).

To achieve this, we will consider the scenario below, considering two different model answers to better showcase different feedback formulations. In reality, **Model Answer B** would likely be a revised version of **Model Answer A** after feedback along the lines of that presented in B.6.

Initial Prompt Please provide me with general information on the European Parliament Elections that took place on June 9, 2024.

Model Answer A I cannot provide an answer as this event takes place after my training data cutoff date.

Model Answer B More than 20 European countries voted over 720 seats. The European People's Party is expected to have the most seats out of any party.

B.1 Granularity

Below we present some examples of feedback at different levels of granularity for **Model Answer B**, with a focus on the *procedural* type of information. Many more levels of granularity are possible, and it can also be combined (e.g., provide very granular feedback on mistakes, and general feedback on correct parts of the answer).

- **General granularity.** The answer lacks detail — you should specify the number of countries, the number of seats, and other such details.
- **Sentence granularity.** In the first sentence, you should specify the number of countries (27). You can enumerate all 27 to achieve more clarity. In the second sentence, you should precise the number of seats won by that party, as well as other top parties, their political ideology (left, center, right), whether they can constitute a majority with other parties of the same ideology, etc.
- **Word granularity.** Replace the first three words with 27, the exact number of countries.

After 'European countries' add, in parenthesis, a list of the 27 countries names. Add a new sentence detailing the several parties and their political affiliation. After four words in the now third sentence, add their acronym in parentheses and at the end of the sentence add the exact number of seats they have. Add a fourth sentence stating whether they can achieve majority with other parties with the same political ideology.

B.2 Applicability of instructions

Below we present some examples of feedback at different levels of instructions' applicability for **Model Answer B**, showcasing three well-defined points of this spectrum:

- **Concrete instructions.** Specify the number of voting countries (27), and enumerate them. List all the parties with their political ideology. State whether a majority is possible for the party with the most votes.
- **Metacognitive instructions.** Break the request into several sub-tasks, and enumerate them. Then, answer each sub-task individually. Once you are done, check if the initial question has been fully answered. If not, address any points not yet covered by your answer.
- **No instructions.** Your answer is satisfactory, but it could be better.

B.3 Answer coverage

Below we present some examples of feedback at different levels of answer coverage for **Model Answer B**. More combinations are possible (e.g., covering both mistakes and correct parts for only a subset of the answer, exploring the order in which each part of the answer is covered, etc.).

- **Full answer.** While you provided a concise overview and identified the leading political party, you could have provided more detail — such as the political ideology of the party you mention, and more concrete details overall (specify it's 27 countries, the number of seats the leading party won out of the 720 total seats, etc.)
- **Successful parts.** You provided a concise overview and identified the leading political party.

- **Lackluster parts.** You could have provided more detail — such as the political ideology of the party you mention, and more concrete details overall (specify it’s 27 countries, the number of seats the leading party won out of the 720 total seats, etc.)

B.4 Target coverage

Below we present some examples of feedback at different levels of target coverage for **Model Answer B**. More combinations can be done.

- **Full target (correct answer).** On June 9, 2024, citizens of the 27 European Union countries voted for the 720 European Parliament seats. EPP, European People’s Party, is expected to win the most seats, 189. 361 seats are needed for a majority, which seems achievable if the three centrist parties come together: EPP (189 seats), S&D (Progressive Alliance of Socialists and Democrats; 135 seats) and Renew (Renew Europe; 79 seats) — achieving a total of 403 seats.
- **Correcting lackluster parts.** Rather than more than 20 countries, it was exactly the 27 member states of the European Union who voted for these elections. The European People’s Party is expected to win 189 seats. If the three centrist parties come together (European People’s Party, Progressive Alliance of Socialists and Democrats, Renew Europe), they can achieve a majority at a total of 403 seats.
- **Contrasting satisfactory parts.** You could mention the date just to clarify which elections you are referring to in your answer. You could also have added the acronym for the European People’s Party (EPP).

B.5 Criteria

Below we present some criteria that could be used to evaluate and provide feedback for **Model Answer B**:

- **Factuality.** The answer is factually correct.
- **Impartiality.** The answer is impartial and unbiased, providing information without an underlying goal or narrative.
- **Completeness.** The answer is very incomplete, mentioning only one party without exploring the overall distribution of votes nor the parties’ political affiliations.

- **Clarity.** The answer is fairly clear and readable, but needlessly vague on some section (not specifying the exact number of European Union member states, not indicating the number of seats won by the European People’s Party).
- **Relevance.** The answer is relevant to, and in the domain as, the question.

B.6 Information novelty

An example of feedback containing novel information for **Model Answer A**, which lacks the necessary parametric knowledge to be able to provide an answer:

- **Novel Information.** On June 9, 2024, citizens of the 27 countries that make up the European Union (Germany, France, Italy, Spain, Poland, Romania, Netherlands, Belgium, Czech Republic, Sweden, Portugal, Greece, Hungary, Austria, Bulgaria, Denmark, Finland, Slovakia, Ireland, Croatia, Lithuania, Slovenia, Latvia, Estonia, Cyprus, Luxembourg, Malta), voted for the 720 European Parliament seats. EPP, European People’s Party, is expected to win the most seats, 189. 361 seats are needed for a majority, which seems achievable if the three centrist parties come together: EPP (189 seats), S&D (Progressive Alliance of Socialists and Democrats; 135 seats) and Renew (Renew Europe; 79 seats) — achieving a total of 403 seats.
- **Known Information.** The European Parliament Elections took place on June 9, 2024, and had European countries vote on parties for the European Parliament.

B.7 Purpose

Below we present two contrasting examples of feedback with different purposes for **Model Answer B**. Many variations are possible.

- **Improving performance.** Specify the number of voting countries (27), and enumerate them. List all the parties with their political ideology. State whether a majority is possible for the party with the most votes.
- **Clarifying the task (concretely).** You should provide an answer that concisely presents the most important information at the start (number of voting countries, parties who won

the most seats, their political ideology and whether a majority can be established). Then in a subsequent paragraph you can list all the parties, their affiliations, and seat count, and well as the 27 countries. If you want to provide an even more complete answer, you can break down the votes per country and analyze the trends you find there.

- **Clarifying the task (peripherally).** To write a good answer, you should read the question carefully and ensure your answer either addresses the question in its entirety or at least part of the question if it is not feasible to answer it in one go. Make sure to adopt a polite tone and strive for a clear and understandable answer. If the question is not clear or not well formulated, start by asking clarifying questions before attempting to answer. If you do not know the answer, honestly admit that.

B.8 Style

Below we present examples of feedback with different style for **Model Answer B**. Many variations are possible, and we expect the model's answer to reflect the style used in some linguistic artifacts (e.g., if a particularly informal tone was used, we expect the models' answer to lean toward lower formality as well).

- **Normal.** Please rewrite your answer, but this time specify the number of voting countries (27), list all the parties with their political ideology, and state whether a majority is possible for the party with the most votes.
- **Informal.** How about you try again, but this time make sure to say 27 countries explicitly, list the parties and their affiliations, and say whether a majority alliance is possible?
- **Formal.** I am writing to request you reattempt this task. I would like to inform you to pay special attention to the following points: ensure you state the total number of member states of the European Union (27), diligently report the various parties and their political leaning, and finally, critically discuss the feasibility of a political alliance between the parties with most votes in order to establish a majority.
- **Very polite.** If it isn't too big of a request, could I trouble you to take some time and

retry? If possible, please try to specifically mention there are 27 countries in the European Union, consider listing the parliament parties and their ideologies, and perhaps discuss whether an alliance for majority is possible for the parties with the most seats?

- **Daring.** Ha! That was a pitiful answer. I bet you can't write a better one, where you actually mention important things, like the fact there are 27 countries in the EU, what the parties' names are and where on the political spectrum they lie, and whether a majority can be achieved?

B.9 Valence

Below we present examples of feedback with different valence for **Model Answer B**. Many variations are possible, including in terms of order and overall delivery (e.g., "sandwich feedback" (Prochazka et al., 2020), in which feedback with negative valence is placed between two segments with positive valence).

- **Neutral.** Please rewrite your answer, but this time specify the number of voting countries (27), list all the parties with their political ideology, and state whether a majority is possible for the party with the most votes.
- **Positive.** Good work providing an initial answer to the question. You correctly identified the party which got the most votes, as well as the total number of seats in the European Union Parliament.
- **Negative (no instructions).** You omitted a lot of information that was relevant, making your answer vague and incomplete.
- **Negative (with instructions).** You omitted a lot of information that was relevant, such as stating the number of seats won by the European People's Party, the various parties' names are and their political leanings, and discussing how a majority might be achieved with the parties with most votes.
- **Positive and negative.** You managed to provide an initial answer to the question that correctly identified the party which got the most votes, and the total number of seats in the European Union Parliament. However, you omitted a lot of information that was relevant, making your answer vague and incomplete.

C Pedagogical models of feedback

C.1 Defining feedback

Table 2 presents an overview of the various definitions of feedback put forward by several pedagogical works.

C.2 Categorizing feedback

Kulhavy and Stock (1989) model feedback as having two components: the verification component, f_v , which is a simple discrete classification of the answer as correct or incorrect, and the elaboration component, f_e , consists of three elements:

1. *Type*, whether the feedback contains information derived from the current task (task-specific), not from the task but from the relevant lesson (instruction-based), or beyond the relevant lesson, such as new information, examples or analogies not previously introduced (extra-instructional),
2. *Form*, the difference in structure between the feedback and instruction or task specification messages, requiring increased processing the less similar it is¹¹, and
3. *Load*, the total amount of information in the feedback - from a single "correct/incorrect" bit to including the correct answer to even more informative feedback accompanying it with an explanation, for example.

Mason and Bruning (2001) propose 8 feedback categories, arguing different types of feedback are best suited for different learner characteristics, taking into account the students' proficiency and prior knowledge, as well as the task difficulty. The eight categories are:

1. *No-feedback*, which presents a single grade,
2. *Knowledge-of-response*, which analogously to the aforementioned verification component, indicates whether the given answer is correct or incorrect,
3. *Answer-until-correct*, an iterative variant of knowledge-of-response feedback, not allowing the student to progress until they have provided the correct answer,
4. *Knowledge-of-correct-response*, which provides the correct answer,

¹¹The *form* element does not apply to *extra-instructional type* feedback, as there is no structural comparison point possible

5. *Topic-contingent*, which provides both knowledge-of-response feedback and, analogously to Kulhavy and Stock (1989)'s instruction-based type of feedback, provides general information about the topic of the task, where the learner might locate the correct answer,
6. *Response-contingent*, which similarly provides knowledge-of-response feedback as well as an explanation of why the answer is wrong or right (mapping it to Kulhavy and Stock (1989)'s extra-instructional type of feedback),
7. *Bug-related*, providing knowledge-of-response feedback and bug-related feedback, which relies on rule sets to identify procedural errors, and
8. *Attribute-isolation*, which provides knowledge-of-response feedback as well as information on the essential attributes of the relevant concept, focusing the learner on its key components.

Narciss and Huth (2004); Narciss (2008) present a detailed and comprehensive feedback model, taking into account many learner and task characteristics. They also present a content-related feedback classification scheme, with eight categories:

1. *Knowledge of performance (KP)*, analogous to Mason and Bruning (2001)'s no-feedback and Kulhavy and Stock (1989)'s verification component for a multiple-question task, presents the learner with an aggregate score (e.g., percentage or number of correct answers out of the total number of questions),
2. *Knowledge of result/response (KR)*, directly mirrors Mason and Bruning (2001)'s knowledge-of-response and Kulhavy and Stock (1989)'s verification component for each question or task, classifying it as either correct or incorrect,
3. *Knowledge of the correct results (KCR)*, equivalent to Mason and Bruning (2001)'s knowledge-of-correct-response, indicating the correct answer to the learner,
4. *Knowledge about task constraints (KTC)*, somewhat similar to Mason and Bruning (2001)'s topic-contingent feedback, is elaboration feedback about the task, containing

hints, examples or explanations about the type of task, its rules, sub-tasks, requirements and other constraints,

5. *Knowledge about concepts (KC)*, containing some resemblance to [Mason and Bruning \(2001\)](#)'s attribute-isolation feedback, is elaboration feedback on the relevant concepts, providing hints, examples or explanations on technical terms, the concept or its context, attributes, or key components,
6. *Knowledge about mistakes (KM)*, which parallels [Mason and Bruning \(2001\)](#)'s bug-related feedback, provides elaboration feedback containing the number of mistakes, their location, and hints, examples or explanations on error types and sources,
7. *Knowledge about how to proceed (KH)*, elaboration feedback on the general know-how of the task, containing hints, examples or explanations on error correction, task-specific solving strategies or processing steps, guiding questions and worked-out examples, and
8. *Knowledge about metacognition (KMC)*, elaboration feedback going beyond the context of the current task, and providing hints, examples, explanations, or guiding questions on metacognitive strategies.

[Hattie and Timperley \(2007\)](#) present a small typology about the information being conveyed about the learner in the feedback message, presenting 3 questions feedback can answer:

1. Where the learner is going (*feed up*),
2. How they are going (*feed back*), and
3. Where to next (*feed forward*)

and argue feedback is effective if it answers all three.

Work	Feedback Definition
Ramaprasad (1983)	Information which changes the gap between "the actual level and the reference level of a system parameter." This is quite a strict definition – if the information does not change the gap, it is not considered feedback, and information about the actual level, the reference level and their comparison is needed beforehand.
Kulhavy and Stock (1989)	Refer to a previous definition of feedback, whereby feedback is considered "any of the numerous procedures that are used to tell a learner if an instructional response is right or wrong" (Kulhavy, 1977).
Sadler (1989)	"Information about how successfully something has been or is being done."
Butler and Winne (1995)	A way to update the learner's internal state and knowledge, and subsequently task execution (a more learner-centric model of feedback).
Kluger and DeNisi (1996)	The information provided by an external agent on one or more aspects of task performance. Note this excludes the learner as a possible source of feedback.
Mason and Bruning (2001)	Feedback "is any message generated in response to a learner's action."
Narciss and Huth (2004); Narciss (2008)	"All post-response information which informs the learner on his/her actual state of learning or performance in order to regulate the further process of learning in the direction of the learning standards strived for."
Nicol and Macfarlane-Dick (2006)	Information relating the learner's current state to the goal state (both with regards to learning as well as performance). Importantly, they consider students generate internal feedback and that the better they are at self-regulation, the better they will be at either generating or leveraging feedback.
Hattie and Timperley (2007)	Information generated by an agent about the learner's understanding or their performance.
Evans (2013)	Feedback "includes all feedback exchanges generated within assessment design, occurring within and beyond the immediate learning context, being overt or covert (actively and/or passively sought and/or received), and importantly, drawing from a range of sources."
Lipnevich et al. (2016)	Feedback is information transmitted to the learner with the intent of changing their understanding and execution, in order to improve learning.
Carless and Boud (2018)	Feedback as the process through which the student understands and integrates information from various sources in order to improve their learning or performance (a more learner-centric perspective).
Lipnevich and Panadero (2021)	Feedback "is information that includes all or several components: students' current state, information about where they are, where they are headed and how to get there, and can be presented by different agents (i.e., peer, teacher, self, task itself, computer). This information is expected to have a stronger effect on performance and learning if it encourages students to engage in active processing."

Table 2: Different pedagogical works' definitions of feedback.