

Navigating the Landscape of Hint Generation Research: From the Past to the Future

Anubhav Jangra^{◇*} Jamshid Mozafari[†] Adam Jatowt[†] Smaranda Muresan[◇]

[◇]Department of Computer Science Columbia University, USA

[†]Department of Computer Science University of Innsbruck, Austria

Abstract

Digital education has gained popularity in the last decade, especially after the COVID-19 pandemic. With the improving capabilities of large language models to reason and communicate with users, envisioning intelligent tutoring systems that can facilitate self-learning is not very far-fetched. One integral component to fulfill this vision is the ability to give accurate and effective feedback via hints to scaffold the learning process. In this survey article, we present a comprehensive review of prior research on hint generation, aiming to bridge the gap between research in education and cognitive science, and research in AI and Natural Language Processing. Informed by our findings, we propose a formal definition of the hint generation task, and discuss the roadmap of building an effective hint generation system aligned with the formal definition, including open challenges, future directions and ethical considerations.

1 Introduction

Prior research has established a correlation between the student-teacher ratio and a student's overall performance (Koc and Celik, 2015). However, private tutoring is not accessible to everyone, and finding expert tutors is often difficult and incurs considerable costs (Bray, 1999; Graesser et al., 2012). Intelligent tutoring systems (ITSs) hold the key to addressing these educational challenges, notably the need for personalized learning in a system often reliant on instructional teaching and standardized testing (Anderson et al., 1985).

The hallmark of intelligent tutoring systems is their ability to provide step-by-step guidance to students while they work on problems, and hints play a critical role in their ability to provide this

help (Biswas et al., 2014). Hints are a tool to provide scaffolded support to the learners, and can be traced back to the socio-cultural theory of Vygotsky's Zone of Proximal Development, referring to "*the gap between what a learner can do without assistance and what a learner can do with adult guidance or in collaboration with more capable peers*" (Vygotsky, 1978).

Within learning science, hints refer to the clues, prompts, questions, or suggestions provided to learners to aid them in solving problems, answering questions, or completing tasks, thereby encouraging critical thinking, problem-solving skills, and independent learning. In Figure 1, we provide an example of a hint generation system capable of the reasoning required to answer the question, acknowledging the wrong attempt by the learner and providing informative hints linked to the learner's existing knowledge. This framework of scaffolding is well established in education (Van de Pol et al., 2010), and expert tutors are guided to incorporate it in their teaching practices (Belland, 2017).

With the aim of developing hint generation systems with capabilities such as the ones showcased in Figure 1, we consolidate the dispersed efforts on hint generation, bridging the gap between research in education and cognitive sciences on the one hand, and research in AI and Natural Language Processing (NLP) on the other. Grounded in the findings from our literature review, we provide a roadmap for future research on automatic hint generation. We summarize the key characteristics of a successful hint as observed by research with human tutors in Section 2 and review the automated hint generation systems in Section 3. We identify the gaps and propose a roadmap for future research in hint generation in Section 4. We provide a rethinking of the formal definition of the

* Corresponding author - anubhav@cs.columbia.edu.

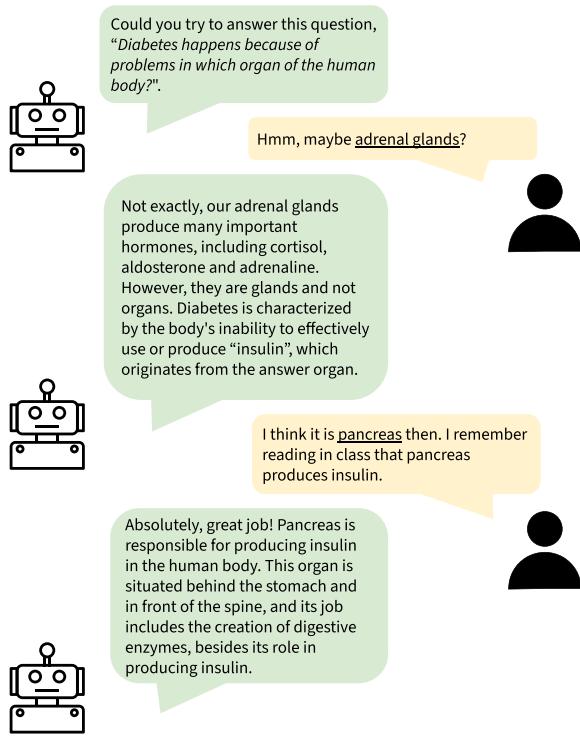


Figure 1: An example of a hint generation system capable of acknowledging the learner’s wrong answer, and scaffolding them to the correct direction.

hint generation task (Section 4.1), a brief review of research areas that can inform the design of a hint generation system that aligns with the formal definition (Section 4.2), open challenges for future directions for effective automatic hint generation systems (Section 4.3), and ethical considerations (Section 4.4). Our major contributions include:

1. A literature review on hint generation that bridges the gap between research in education and cognitive science on the one hand and research in AI and NLP on the other.
2. A formal definition of the hint generation task, grounded in the cognitive theories on learning and findings from qualitative research.
3. A roadmap for research on automatic hint generation, outlining challenges, promising future directions, and ethical considerations for the field.

2 Anatomy of a Hint

In this section, we draw on research from education and cognitive sciences to describe the key characteristics of an effective hint formulation

process. We start by describing the pragmatics of a hint ('context'), covering some prominent traits exhibited by expert tutors and educators while generating hints, and then dive deeper into the anatomy of a hint by discussing the *semantic* ('what to say') and the *stylistic* ('how to say it') aspects of a hint.

2.1 Pragmatics of a Hint

Expert tutors are able to provide high quality support to students because they are aware of each student’s individual learning style, strengths, and weaknesses. These tutors often exhibit a contextual awareness about their students, and we describe two such common practices adopted by educators’ when supporting students below.

Scaffolding Support: Structuring hints in a scaffolded manner, with incremental steps leading to the solution, helps learners systematically build their understanding. These “just-in-time interventions” (Wood et al., 1976) allow students to build their understanding step by step, starting with foundational concepts and progressing toward more advanced aspects without being overwhelmed by the information or task complexity (Zurek et al., 2014; Lin et al., 2012; Hammond, 2001).

Personalization and Learner Feedback: Every learner is unique and has different needs and preferences when it comes to learning (Bulger, 2016). Prior studies point towards a learner-centered pedagogical system, where personalization and individualization of learning have a significant role in the students’ overall learning process and strengthen their sense of self and individuality (Radovic-Markovic and Markovic, 2012). To generate effective hints, it is important to recognize and cater to these individual needs by considering learners’ strengths, challenges, cultural sensitivity, and preferences (Chamberlain, 2005; Ibrahim and Hussein, 2016; Suaib, 2017). A good learning environment also incorporates a feedback loop, where hints are accompanied by opportunities for learners to provide feedback, promoting active engagement. This two-way communication allows tutors to gauge the effectiveness of their guidance and adjust their learning plans (Boud and Molloy, 2013).

2.2 Semantics of a Hint

Semantics of a hint refers to the information conveyed by the hint, which includes explaining the

key concepts and ideas required to scaffold the learning process. We observed the following properties of effective hint's semantics.

Relevance to the Learning Objective: Learning objectives serve as a measure of achievable goals that articulate what learners should know or be able to do by the end of a learning experience. Learning objectives can broadly be categorized across three domains: cognitive, affective, and psychomotor objectives (Hoque, 2016; Sönmez, 2017). Each domain has different expectations and goals to assess the effectiveness of a hint. A hint generation system should model these objectives to create successful high-quality hints.

Link to Prior Knowledge: A successful hint would act as a bridge between a learner's existing knowledge base and the current learning step to foster continuity in learning. Studies have shown that building on prior knowledge helps students bridge gaps, clear misconceptions, and reinforces the relevance of new information (Hailikari et al., 2008, 2007; Dong et al., 2020).

Conceptual Depth: Many learning sessions focus on teaching learners how to harness latent cognitive abilities and mold them into *deep conceptual thinkers* with the ability to discuss and question more, seeking to understand rather than only memorize (Rillero, 2016). It is important to balance the complexity of a hint that strikes a student's interest without overwhelming them.

2.3 Style of a Hint

Expert human tutors adopt diverse techniques to convey information to learners. These strategies vary from non-verbal cues such as body language, facial expressions, and vocal tone (Bambaeeroo and Shokrpour, 2017; Wahyuni, 2018) to adopting multimedia content to teach complex topics (e.g., using animations and maps to teach the geographical concept of “folded mountains”) (Kapi et al., 2017). We try to cover the most relevant aesthetic aspects of hints that might be useful in building better hint-generation system.

Clarity and Simplicity: Hints should be expressed in clear and simple language to ensure that learners easily grasp the underlying concept or problem-solving strategy. Avoiding unnecessary complexity enhances the usefulness of the hint and is usually well received by the learners. This is a well-established practice within the learning sciences community known as *direct in-*

structions (Kozloff et al., 1999; Kim and Axelrod, 2005; Rosenshine, 2008).

Encouragement and Positive Tone: The role of encouragement and positive attitude has been extensively investigated in several human studies in classroom settings, and all unanimously align with the significance of motivating learners towards better performance, increased participation, and improved self-confidence (Ducca, 2014; Yuan et al., 2019; Li, 2021; Lalić, 2005). A hint generation systems could benefit by incorporating a positive, encouraging tone (as demonstrated in Figure 1).

Adopting Creative and Multi-modal Elements: In order to encourage active participation and retain learners' interest, human tutors often adopt several creative and multi-modal elements to facilitate better understanding and information retention. These creative elements include interactive literary devices like analogy (Richland and Simms, 2015; Gray and Holyoak, 2021; Nichter and Nichter, 2003; Thagard, 1992), questions (Hume et al., 1996; Chi, 1996), and metaphors (Low, 2008; Sfard, 2012; Guilherme and Souza de Freitas, 2018). We can also expand beyond text, and incorporate information from other modalities such as maps (Winn, 1991), diagrams (Winn, 1991; Swidan and Naftaliev, 2019; Tippett, 2016), and multimedia content (Abdulrahaman et al., 2020; Collins et al., 2002; Kapi et al., 2017) to effectively complement the learning experience. A good hint can take inspiration from some of these creative elements for a successful transfer of knowledge.

A good instructor typically takes the general guidelines into consideration and uses a mixture of the aforementioned semantic and stylistic features to create effective hints based on their prior tutoring experiences. For instance, to develop a hint that uses the literary device of analogy, the tutor must understand the prior knowledge of the learner to create successful hints (Gray and Holyoak, 2021).

3 Survey of Computational Approaches

In this section, we provide a comprehensive overview of the recent advancements of computational approaches for automatic hint generation. We first describe the extensively studied hint generation techniques for computer programming that focus on revealing code snippets to help learn

Data Source	Input	Expected or Generated outputs
ITAP	<u>Question-</u> Write a program that if a given day is weekend. <u>Learner's response-</u> <pre>def isWeekend(day): return bool(day=='sunday' or day =='saturday')</pre>	Hint generated by Rivers and Koedinger (2017)- Type: Replace Old expression: "saturday" New expression: "Saturday"
ReMath	<u>Question-</u> Mike has 4 cookies and he eats 3 cookies. So Mike has ____ cookies left? <u>Learner's response -</u> He has 10 cookies left.	Expected output Error type- guess Response strategy- provide a solution strategy Response intention- <i>Help student understand the lessons topic or solution strategy</i> Response- <i>Great try! Let's try to draw a picture.</i> <i>Let's start with 4 cookies and erase the 3 that Mike eats.</i>
SQuAD	<u>Question-</u> Who became the most respected entrepreneur in the world according to Financial Times in 2003? <u>Expected Answer-</u> Bill Gates	Hint Generated by Jatowt et al. (2023)- <i>The searched person held the position of chief executive officer.</i>
TrivialQA	<u>Question-</u> In which city are the headquarters of the International Monetary Fund? <u>Expected Answer-</u> Washington D.C.	Hint Generated by Mozafari et al. (2024b)- <i>The city is known for its neoclassical architecture.</i> <i>The city is located on the Potomac River.</i> <i>The city is the capital of the USA located on the east coast.</i>

Table 1: Hint generation examples from selected prior research discussed in Section 3.

how to program. Next, we dive into the relatively under-explored natural language hint generation, where we explore strategies for diverse domains like mathematics, language acquisition, or factual entity-based questions. We conclude this section by describing some limitations of automatic hint generation systems today, and propose a roadmap for future research in the field in Section 4.

3.1 Hint Generation for Computer Programming

A vast majority of computational approaches for hint generation have focused on the specific domain of computer programming, owing to the more objective nature of the task and abundance of data. We briefly discuss the approaches, datasets, and evaluation metrics adopted in the field. For a more comprehensive review in this specific domain, we refer the readers to the surveys written by Le et al. (2013), Crow et al. (2018), McBroom et al. (2021), and Mahdaoui et al. (2022).

Datasets. Two widely popular datasets in the programming hint generation space are *iSnap* (Price et al., 2017) and *ITAP* (Rivers and Koedinger, 2017). Both datasets consist of detailed logs collected from several students working on multiple programming tasks, including the complete traces of the code and records of when the hints were requested. *iSnap* (Price et al., 2017) is based on *Snap!*¹—a block-based educational

graphical programming language, while *ITAP* (Rivers and Koedinger, 2017) is a Python dataset collected from two introductory programming courses taught at Carnegie Mellon University. In Table 1, we describe an example from the *ITAP* dataset, where the goal is to write a program to determine if a given day is weekend. Given a student’s code that fails to pass the pre-determined test cases (e.g., `isWeekend("Sunday")` will return `False`), the aim of a hint generation system is to provide hints to help them successfully solve the problem (e.g., replacing lowercase ‘saturday’ to uppercase ‘Saturday’).

Approaches. Most of the recent efforts in programming hint generation adopt a data-driven deterministic approach (Barnes and Stamper, 2008; Rivers and Koedinger, 2017; Obermüller et al., 2021; Jin et al., 2012; Zimmerman and Rupakheti, 2015; Paaßen et al., 2018; Price et al., 2016; Rolim et al., 2017), that comprises three key components: a corpus of diverse candidate solutions (usually obtained via past student attempts), a matching algorithm to select the best candidate response given an ongoing attempt based on similarity, and graph-based solution path construction to synthesize hints (Figure 2 describes one approach in detail). We found abstract syntax tree (ASTs) (McCarthy, 1964; Knuth, 1968) to be the most popular choice of graph representation for hint synthesis due to its vast literature and language-agnostic nature. Similarly, McBroom

¹<https://snap.berkeley.edu/>.

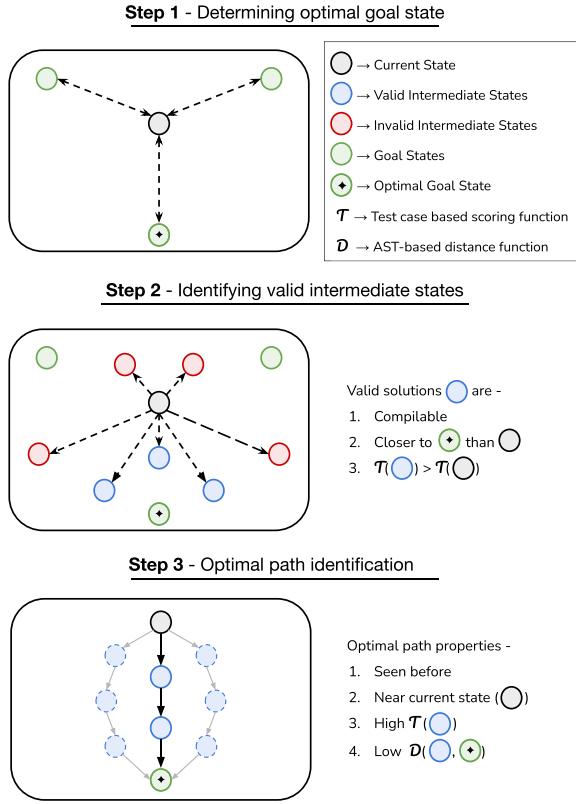


Figure 2: Illustration of the *path construction algorithm* (Rivers and Koedinger, 2014), which generates programming hints for an ongoing student attempt (current state), given reference solution(s) (goal states) and a test case-based scoring function T .

et al. (2021) provide a detailed generalization of hint generation techniques for the programming domain called HINTS (Hint Iteration by Narrow-down and Transformation Steps framework). Although effective, programming hints are rarely natural language responses and, therefore, are not capable of incorporating the stylistic aspects of hints that improve the learner’s experience (Section 2.3). We believe an amalgamation of NLP technologies and current hint generation systems could improve upon this limitation.

Evaluation Metrics. Given the vast majority of work adopting a data-driven hint strategy, programming hints are designed to reveal some aspect of the program not written by the learner. Within our survey, all the evaluation metrics adopted a code-based similarity measure to gauge the quality of the hint, and we found two prominent evaluation paradigms that we categorize into *reference-based* and *reference-free* evaluation metrics. The *reference-based* evaluation metrics (Price et al., 2019) assume the availability

of reference hints developed by expert tutors. For instance, *QualityScore* (Price et al., 2019) is a reference-based evaluation metric that uses abstract syntax tree-based similarity measure to evaluate the quality of the generated programming hints with respect to expert tutor written hints.

Reference-free evaluation metrics (Rivers and Koedinger, 2017; Paaßen et al., 2018; Obermüller et al., 2021; Zimmerman and Rupakheti, 2015) focus on measuring the impact of a hint to the learner’s response. For example, Paaßen et al. (2018) proposed root-mean-square error over two distance measures: (1) distance between the predicted post-hint state and true next state, and (2) distance between the predicted post-hint state and learner’s true final state. Rivers and Koedinger (2017), on the other hand, adopted a learner-agnostic approach to evaluate the hint quality. They used two measures of a successful chain of hints: (1) the ability of the hint sequence to reach the correct solution state and (2) the length of the chain.

Reference-based evaluation metrics are able to compare the quality of generated hints with the expert tutor’s feedback capabilities, but can be difficult to scale. *Reference-free* evaluation metrics, on the other hand, are capable of evaluating previously unseen problems but require human evaluations to obtain performance signals for the metrics. Each of these paradigms has pros and cons, but currently, there is no holistic measure of the quality of programming hints that emphasizes the pragmatics, semantics and style characteristics of the generated hints.

3.2 Natural Language Hint Generation

Natural language hint generation has risen to popularity in the last few years as a consequence of the recent advancements in NLP, particularly large language models capable to generate fluent and coherent text (Min et al., 2023; Naveed et al., 2023). We found the question answering format to be the most prevalent setup for natural language hint generation systems, where the learner attempts to answer a question to recall and concertize their understanding of a concept.

Datasets. ReMath (Wang et al., 2023) is a benchmark co-developed with math teachers (*i.e.*, experts) for evaluating and tutoring students in

the mathematics domain. ReMath provides a systematic breakdown of the human-tutoring process into three steps: (1) identifying the error type, (2) determining a response strategy and intention, and (3) generating a feedback response that adheres to this strategy (example in Table 1). Each of these steps is manually annotated by an expert math teacher. Wang et al. (2023) also provide a set of error types (*e.g.*, guess, careless, misinterpret, right-idea), response strategies (*e.g.*, explain a concept, ask a question), and intentions (*e.g.*, motivate the student, get the student to elaborate the answer) to facilitate the feedback generation process. ReMath is a great example of high-quality data collection with human experts towards building better hint generation systems.

TriviaHG (Mozafari et al., 2024b) is another hint generation dataset developed by extending the TriviaQA dataset (Joshi et al., 2017). Mozafari et al. (2024b) utilize Microsoft’s CoPilot² to generate hints due to its retrieval augmented generation approach that generated more reliable responses grounded on internet-retrieved documents. TriviaHG includes 10 generated hints for questions that CoPilot is capable of answering. In a follow-up work, Mozafari et al. (2024a) propose HintQA that utilizes these hints as concise context, improving the QA capabilities of LLMs over other context-retrieval and context-generation baselines.

Approaches. For open-ended hint generation of factoid questions, Jatowt et al. (2023) proposed a Wikipedia-based retrieval framework for the “Who?”, “Where?” and “When?” type questions. They propose a popularity-based framework for the “When” question type, where the popularity of an event for the answer year is measured by the count of Wikipedia hyperlinks directing to the event’s website, and a hand-curated template approach for the “Who” and “Where” question types (example hint described in Table 1). On the other hand, Wang et al. (2023) benchmarked the ReMath dataset by instruction fine-tuning (Wei et al., 2021) the language models like Flan-T5 (Chung et al., 2022) and GODEL (Peng et al., 2022), and using in-context learning (Dong et al., 2022) prompts for gpt-3.5-turbo and gpt-4 (Achiam et al., 2023). Tack and Piech (2022) and Wang et al. (2023) found the direct use of

LLMs to fall short in comparison to human expert responses.

Pal Chowdhury et al. (2024) propose a hint generation framework for middle-school level math word problems (termed MWPTutor). MWPTutor provides a hint by formulating a question around the next operation to be performed in the state space. This hint is obtained by matching the ongoing response with all possible decomposed solutions obtained by using a language model to decompose the solutions into atomic mathematical operation steps.

Current question answering hint generation systems do not personalize the hints to learners’ preferences, learning objectives, or their prior knowledge (Sections 2.1 and 2.2). We discuss how we can improve these hint generation systems to aid the learning process in Section 4.

Evaluation Metrics. All the discussed approaches for hint generation have used human evaluation to assess the quality of the system’s output. Jatowt et al. (2023) conducted a between-subjects study to evaluate their proposed hint generation strategies across different experimental groups. Tack and Piech (2022) proposed the “AI Teacher Test”, comparing the generated responses against the teacher responses across three dimensions—“*speak like a teacher*”, “*understand a student*”, and “*help a student*”. They identified that the LLMs are good at conversation uptake (*i.e.*, the first two requirements) but are quantifiably worse than real teachers on several pedagogical dimensions, especially *helpfulness to a student*. Wang et al. (2023) evaluate the error type identification and feedback response strategy selection as a multi-class classification task and utilize exact match and Cohen’s kappa to measure the accuracy and entropy to measure the output diversity. They also conducted human evaluations for the response generation task and found that all models constrained by knowledge of ground-truth error type and response strategy outperformed their unconstrained counterparts.

For evaluation of hint generation systems in quantitative user studies, Pal Chowdhury et al. (2024) extends the idea of *success rate* (S : fraction of correctly answered questions) and *telling rate* (T : fraction of conversations where the answer was revealed) proposed by Macina et al. (2023). Pal Chowdhury et al. (2024) suggested an *adjusted success* ($S - T$) to prevent an overtly

²<https://copilot.microsoft.com/>.

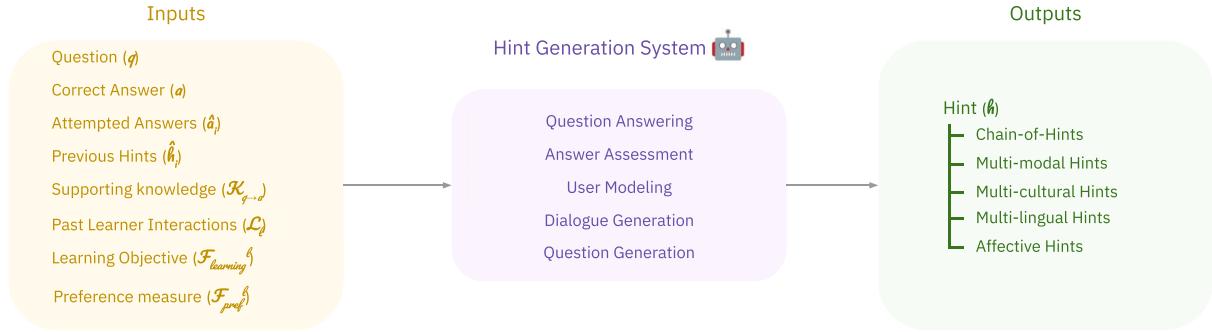


Figure 3: Roadmap of the proposed computational hint generation framework.

revealing framework from achieving high performance, and the harmonic mean of success rate and adjusted success rate as the overall *tutoring score* ($\frac{2S(S-T)}{2S-T}$).

Mozafari et al. (2024b), on the other hand, proposed two learner-agnostic automatic evaluation metrics: *convergence* to measure the ability of a hint to eliminate wrong candidate answers, and *familiarity* to measure the recognizability of answer entities. To measure *convergence* of a hint, they adopt a three-step process: (i) generating candidate answers using LLMs, (ii) validating the entailment of answer given a hint, and (iii) computing an aggregate score for a hint across all candidate answers. For *familiarity*, they utilize the Page Views of Wikipedia³ articles corresponding to the named entities present in the hint as a measure of global familiarity to the hint normalized across all questions in the TrivialHG corpus (Mozafari et al., 2024b).

4 Roadmap for Future Research in Hint Generation

Great progress has been made in automatic hint generation over the last two decades; however, there is still room for improvement. Existing hint generation frameworks do not personalize the hints to the learner’s prior knowledge (Section 2.1), and are only evaluated in short-term studies. We still do not know the effects of long-term exposure to these interventions on the learners. These frameworks are also limited to certain domains and have not been widely explored in other domains, including different branches of science and social science. To improve upon these factors, we discuss a roadmap for future efforts in automatic hint

generation. We propose a computational hint generation framework that draws on education and cognitive sciences (Figure 3).

4.1 Formal Definition

Jatowt et al. (2023) proposed a formal definition for hint generation as follows: Given a question q and its correct answer a , the task is to generate a hint h , such that $P(a|q, h) - P(a|q) > \epsilon$, where $P(a|q, h)$ denotes the probability of a user answering q if the hint h is given, $P(a|q)$ is the probability of a user answering q without the hint, and ϵ is a threshold parameter ($\epsilon > 0$). This definition emphasizes the hint’s ability to help answer a question and does not incorporate any pedagogical aspects or principles. It also does not take individual preferences into consideration.

Below, we provide a comprehensive definition that draws inspiration from widely adopted cognitive frameworks on human learning, such as Anderson’s *adaptive control of thought rational (ACT-R) theory* (Anderson, 2013) and Ausubel’s *theory on meaningful learning* (Ausubel, 1963, 1962, 2012). We also incorporate the findings from the qualitative experiments discussed in Section 2, explicitly integrating the alignment of hints to students’ learning objectives and their prior knowledge (Section 2.2), as well as incorporating the pragmatics of a hint (Section 2.1) by accounting for learners’ preferences. We formulate the hint generation task within a tutoring framework with the goal of correctly answering a question. We will later explain how the definition can be extended to generate hints for tasks beyond question answering.

Refined Formal Definition. Given a learner l attempting to answer a question q , a **hint generation system H** (Figure 3) generates a hint $h \in \mathcal{H}$

³<https://www.wikipedia.org/>.

by mapping $\mathbf{H} : \mathcal{I} \rightarrow \mathcal{H}$, where $\mathcal{I} = \{q, a, \mathcal{K}_{q \rightarrow a}, D_q^l, \mathcal{L}_l, \mathcal{F}_{learning}^l, \mathcal{F}_{pref}^l\}$ is the input to the hint generation system having the following elements:

- q : question
- a : correct answer
- $\mathcal{K}_{q \rightarrow a}$: supporting knowledge for the question answer pair $\langle q, a \rangle$,
- $D_q^l = \{q, \hat{a}_1, \hat{h}_1, \hat{a}_2, \hat{h}_2, \dots, \hat{a}_k\}$: ongoing dialogue, where \hat{a}_i and \hat{h}_i are, respectively, the learner's past attempts and hints related to q ,
- $\mathcal{L}_l = \{D_{q_i}\}$: learner l 's past learning history,
- $\mathcal{F}_{learning}^l$: a function to measure the learner l 's learning objective(s), and
- \mathcal{F}_{pref}^l : a measure of learner l 's preference of and familiarity to a hint.

H generates the hint h that does not contain the answer (Equation (1)), helps the learner to answer the question (Equation (2)), and aligns with the learner's learning objective(s) (Equation (3)).

$$P(a|q, h, D_q^l) < 1 \quad (1)$$

$$P(a|q, h, D_q^l) - P(a|q, D_q^l) > \epsilon_p \quad (2)$$

$$\begin{aligned} \mathcal{F}_{learning}^l(q \rightarrow D_q^l \rightarrow h \rightarrow a) - \\ \mathcal{F}_{learning}^l(q \rightarrow D_q^l \rightarrow a) > \epsilon_f \end{aligned} \quad (3)$$

If the hint generation model **H** is capable of generating multiple hints that satisfy the Equations (1) to (3) such that **H** returns n hints $\{h_1, h_2, \dots, h_n\}$ in decreasing order of preference, then

$$\begin{aligned} (\mathcal{F}_{pref}^l(h_i) > \mathcal{F}_{pref}^l(h_j)) \rightarrow (i < j) \\ \forall h_i, h_j \in \{h_1, h_2, \dots, h_n\} \end{aligned} \quad (4)$$

Both Anderson's ACT-R theory and Ausubel's meaningful learning theory imply the significance of connecting the new knowledge required to solve a problem to the existing concepts in the learner's knowledge base (entailed in Equation (4)). We also account for the diversity of learners' motivation to study and incorporate the notion of improving an individual's learning objective (Equation (3)). The hint generation strategy for

a learner aiming to improve their answer accuracy would greatly differ from someone aiming to maximize the diversity of acquired knowledge in a learning session.

Extending the Formal Definition. The given definition assumes a question-answering setup for the hint generation task, with an additional assumption of the availability of objective answers. We can modify this definition to accommodate for other hint generation settings as described below.

- For subjective questions, we should replace the correct answer a with an evaluation rubric instead $\mathcal{R} : \mathcal{H} \rightarrow \mathbb{R}$.
- For a hint generation system for the writing assistance task, we need to replace the question with the writing task description ($q \rightarrow T$) and ongoing dialogue with an interactive sequence of learner's writings ($\hat{a}_i \rightarrow \hat{w}_i$), and past hints (\hat{h}_i), $D_q^l \rightarrow \mathcal{W}_T^l \mid \mathcal{W}_T^l \Leftrightarrow \{\hat{w}_1, \hat{h}_1, \hat{w}_2, \hat{h}_2, \dots\}$, and the answer with a target rubric for writing evaluation $a \rightarrow \mathcal{R}$.
- For a multi-modal hint generation system, we can assume the atomic instances of the dialogue (namely, q , a , h , \hat{a} , and \hat{h}) are constituted of different modalities depending on the task specifications.

4.2 Components of an Effective Automatic Hint Generation System

In this section, we briefly review some NLP research areas that can inform the design of an automatic hint generation system (Figure 3) that aligns with the proposed formal definition introduced in Section 4.1.

Question Answering. Current hint generation systems are framed as scaffolding tools within a question answering (QA) setup, where learners are assessed on their ability to answer questions on a relevant topic. These systems can benefit from research in question answering in two ways.

Firstly, hint generation models should be equipped with reasoning abilities to answer questions. Anderson et al. (1995) emphasizes the significance of an underlying production-rule model that helps to break down a learning goal into achievable subgoals in their cognitive tutoring framework, and we posit the QA systems today can

form the foundation of this production-rule model. Datasets like *StrategyQA* (Geva et al., 2021) and strategies like *Self-Ask* (Press et al., 2022) and *Socratic Questioning* (Qi et al., 2023) are great illustrations of the goal decomposition ability described in the ACT-R cognitive learning theory framework (Anderson et al., 1995). We can adapt these computational question answering frameworks coupled with the answer assessment module (described below) to learn mapping $\mathbf{H} : \{q, a, \mathcal{K}_{q \rightarrow a}, D_q^l\} \rightarrow \mathcal{H}$.

Secondly, question answering systems are quite diverse and can solve complex questions spread across multiple modalities such as knowledge bases (LAN et al., 2022), tables (Jin et al., 2022), images (Srivastava et al., 2021; de Faria et al., 2023), and videos (Zhong et al., 2022), as well as explain the reasoning for the answer (Danilevsky et al., 2020; Schwalbe and Finzel, 2023). These question answering models can form the basis of hint generation systems that take into account incomplete solutions and the necessary reasoning steps to answer the question.

Answer Assessment. Answer assessment plays a crucial role in evaluating and understanding learner responses, enabling the system to provide targeted feedback and adaptive guidance. Accurate answer assessment is foundational for standardizing the grading process, identifying misconceptions, tracking individual progress, and tailoring subsequent hints to address specific learning needs. We can utilize an answer assessment module to identify the missing (or wrong) components of an ongoing answer (\hat{a}_k) juxtaposed to a reference answer (a) and provide appropriate hints to scaffold them towards a correct solution. We can also utilize the answer assessment module to aggregate the mistakes made by a learner over time to track their progress.

Prior work includes several answer evaluation strategies, varying from naive exact match and token overlap approaches to BERT-based semantic strategies for short form answers (Bulian et al., 2022), explainable systems trained on LLM-distilled rationales (Li et al., 2023a), semantic grouping based systems for batch grading (Chang et al., 2022) and multi-modal assessment frameworks to evaluate oral presentations (Liu et al., 2020). For a detailed overview of answer assessment strategies, we point the readers to the survey written by Das et al. (2021).

User Modeling. User modeling is extensively explored within the recommendation systems research area, with the aim of adapting and customizing a service to a user’s specific needs. User modeling holds the key to identifying the learner’s interests and preferences within the hint generation framework (Equation (4)), helping align the hints to the learner’s prior knowledge (Section 2.2) and personalize the feedback (Section 2.1). Using a learner’s past interaction data (\mathcal{L}_l), we can learn a preference function (\mathcal{F}_{pref}^l) based on the success of previous hints in helping the learner achieve their learning goal. We can measure this learning goal ($\mathcal{F}_{learning}^l$) using direct or indirect learner feedback from the user interaction. For instance, we can obtain learners’ self-reported data about the effectiveness of a hint, or develop alternate indicators of success by observing the learner’s behavior across several dimensions, such as the number of interactions, response initiation time, degree of change in response, and frequency of hints asked.

Liu et al. (2022), for instance, utilizes long short-term memory (LSTM) cells (Hochreiter and Schmidhuber, 1997) for knowledge estimation (Corbett and Anderson, 1994) of student’s current understanding based on their past responses for open-ended program synthesis. They utilize these time-varying student knowledge states to predict students’ responses to programming problems, monitoring and analyzing their progress. We can adopt similar user modeling techniques to measure progress, adapt the hints to the learner’s preferences, and design a curriculum (or suggest areas) for further improvement. We direct the readers to the comprehensive survey authored by He et al. (2023) for an extensive overview on user modeling.

Dialogue Generation. Intelligent tutoring systems are mediated by dialogue systems (or conversational agents), providing an interface between ITS and learners. Recent developments in retrieval augmented generation (Gao et al., 2023; Cai et al., 2022; Li et al., 2022), emotion-aware dialogue systems (Ma et al., 2020), and enhanced language understanding capabilities in multi-turn dialogues (Zhang and Zhao, 2021) can help improve the hint generation models, empowering them with abilities to better understand the learner’s queries, understand and respond to complex domains, generate affective responses, and

keep track of long conversation sessions for adaptive tutoring that would help link the generated hints to learner’s prior knowledge (Section 2.2). For further exploration, we invite readers to read the surveys written by Ni et al. (2023), Deriu et al. (2021), and Ma et al. (2020).

Question Generation. Questions play a pivotal role in education, serving to help recall knowledge, test comprehension, and foster critical thinking. Al Faraby et al. (2023) classify neural question generation for educational purposes into three broad categories: (1) Question Generation from Reading Materials, (2) Word Problem Generation, and (3) Conversation Question Generation. We posit a similar use case for question generation to generate hints within conversations to clarify not just the learner’s response but also their understanding and ability to present their ideas. We point the readers to surveys written by Kurdi et al. (2020), Zhang et al. (2021), Das et al. (2021), and Pan et al. (2019) for a comprehensive overview.

Modular Structure of Proposed Approach. Although an ideal hint generation system would be able to satisfy all the criterion described in our formal definition (Equations (1) to (4))—providing hints that are personalized to individual learners preferences and learning objectives—we acknowledge the limitations of technology at the time of this survey (Section 4.3). Keeping that in mind, our proposed formal definition (Section 4.1) and components of hint generation system (Section 4.2) suggests a modular structure, comprising of some essential components (such as Equations (1) and (2) that form the crux of a hint generation system, or the question answering module that helps develop the reasoning to develop the hints) and some complementary modules (like Equations (3) and (4) that helps personalize the hints to individual learners or question generation module that adapts particular hint giving strategies) to further enhance the quality of generated hints. This paper proposes thus a long term roadmap for automatic hint generation systems that can showcase intermediate progress along the way.

4.3 Open Challenges for Effective Automatic Hint Generation Systems

In this section, we outline challenges and future directions for building more effective automatic hint generation systems that align with our formal definition.

Privacy-preserving Self-evolving Frameworks. Current hint generation frameworks limit their applications to a fixed dataset or a pre-defined set of problems, which does not guarantee high quality performance in real-world applications. In order to make these frameworks more effective for the learners, we need to adopt self-evolving frameworks that can incorporate the learner’s feedback and implicit preferences. These self-evolving frameworks should have the capability of identifying user feedback from an ongoing conversation (D_q^l), gathering implicit preferences (\mathcal{F}_{pref}^l) from past learning interactions (\mathcal{L}_l), and adapting their hint generation process to an individual learner while respecting their right to privacy.

Although there is prior work on incorporating human-feedback to improve generation quality (*e.g.*, prompt optimization (Chang et al., 2024; Sahoo et al., 2024), reinforcement learning with human feedback (Kaufmann et al., 2023; Casper et al., 2023)), incorporating these user-modeling aspects in a privacy-preserving manner remains an active area of research (Miranda et al., 2024). Differential privacy (Zhao and Chen, 2022; Yang et al., 2023) and federated learning (Li et al., 2021) could be relevant research avenues to help create these self-evolving frameworks with a responsible approach. For an in-depth review of the state of preserving privacy in LLMs, readers are encouraged to consult the survey by Miranda et al. (2024).

Diverse Domain Exploration. Most efforts within the hint generation space are limited to programming, language acquisition or mathematics, similar to the domain trends in intelligent tutoring systems (Mousavinasab et al., 2021). However, education and tutoring often involve other subjects within natural sciences (*e.g.*, physics, chemistry, biology, earth sciences), social sciences (*e.g.*, history, civics, geography, law), and learning beyond educational institutions (*e.g.*, educating patients about their health conditions effectively (Gupta et al., 2020)). Tackling the hint generation problem in these subjects raises several challenges, such as the evaluation of subjective long-form answers, the need for domain knowledge and enhanced reasoning capabilities (Lyu et al., 2024; Huang and Chang, 2022; Zhang et al., 2023). However, it simultaneously opens up new opportunities to expand the capabilities of hint generation systems, including the development of smaller models that combine the power of pre-trained language

models with the power of adapted knowledge models (e.g., COMET (Bosselut et al., 2019)).

Efforts in question answering can also help provide seed datasets and data annotation strategies to expand the hint generation resources. The literature is rich in education-related datasets ranging from generic datasets like RACE (Lai et al., 2017), LearningQ (Chen et al., 2018), TQA (Kembhavi et al., 2017), and so forth to subject-specific datasets like SciQ (Welbl et al., 2017), ScienceQA (Lu et al., 2022), SituatedQA (Zhang and Choi, 2021), GSM8K (Cobbe et al., 2021), BioASQ (Tsatsaronis et al., 2015), CORD-19 (Wang et al., 2020), and PubMedQA (Jin et al., 2019). To extend these datasets we envision a collaborative approach between the ML practitioners and expert tutors to explore the utilization of hints in classroom or one-on-one tutoring setting across different domains.

Multi-lingual and Multi-cultural Aspects. Prior studies have found a correlation between linguistic and cultural diversity, and capabilities for innovation (Hofstra et al., 2020; Evans and Levinson, 2009). However, we found that the majority of research at the intersection of natural language processing and learning sciences is limited to the English language, either as the mode of education or a subject for language acquisition. Although utilizing multi-lingual large language models (mLLMs) for hint generation can potentially help incorporate the linguistic aspects of the learning experience, providing culturally aware hints still remains a challenging task. Liu et al. (2024) reveals that while mLLMs are aware of cultural proverbs, they struggle to reason with figurative proverbs and sayings. Building benchmarks to evaluate cultural awareness of generation models such as MAPS (Liu et al., 2024), PARIKSHA (Liu et al., 2024), CultureAtlas (Fung et al., 2024), and CUNIT (Li et al., 2024) is an essential first step towards identifying the shortcomings of hint generation frameworks. Moving forward, we need to explore the hint generation capabilities of mLLMs and develop culturally aware systems that accommodate the education for non-English learners, and sustain the diverse cultural and linguistic heritage (Bernard, 1992; Soto et al., 1999).

Multi-modal Elements. Prior qualitative research has established gains from complementing

education with additional modalities like maps, diagrams, and multi-media content (Winn, 1991; Swidan and Naftaliev, 2019; Tippett, 2016; Kapi et al., 2017; Abdulrahaman et al., 2020). Research in intelligent tutoring systems has also explored incorporating certain gamifying elements like badges, leaderboards, narratives, and virtual currency to keep the learners more engaged and motivated (González et al., 2014; Ramadhan et al., 2024). We believe incorporating these cues into a hint generation system would improve the students' memory, understanding, and overall learning experience to help create a holistic educational tool (Pourkamali et al., 2021).

Affective Systems. Affective aspects are often neglected when building intelligent tutoring systems and hint generation systems (Hasan et al., 2020). However, incorporating them into the education pipeline is key for personality development, encouragement and improving self-motivation (Jiménez et al., 2018). Jiménez et al. (2018) show that using affective feedback has a positive impact on students facing learning challenges. Thus, affective hint generation systems that take into consideration the emotional state of the learners when providing feedback is an important area of study.

In recent years, there has been increased interest in building dialogue agents that can adapt to the emotional state of the user in an ongoing dialogue setting. Haydarov et al. (2023) develop a large-scale multi-modal benchmark for visually grounded emotional reasoning-based conversations that use visual stimuli to stir a conversation to test out the emotional reasoning capabilities of multi-modal systems. Li et al. (2023b), on the other hand, proposed a future emotion state prediction framework in spoken dialogue systems to predict the future affective reactions of users based on the ongoing conversation. Resources and findings from the ongoing research in affective dialogue systems could be leveraged and advanced to develop more adaptive and emotion-aware tutoring systems. We point the readers to Ma et al. (2020), Raamkumar and Yang (2022), and Zhang et al. (2024) for a comprehensive review of affective dialogue systems.

Accessible Systems. Another fruitful and challenging area of research is to develop accessible hint generation systems. For people with neurodevelopmental disorders (e.g., attention-deficit/

hyperactivity disorder and autism spectrum disorder and learning disabilities (e.g., dyslexia and dyscalculia), one can modify the hint generation systems to have: (1) *simplified text*: using plain language, and avoid using jargon and complex terminology (Štajner, 2021), (2) *multi-sensory supports*: leveraging a combination of visual, auditory, and kinesthetic modalities to present hints in multiple formats (Vezzoli et al., 2017; Gori and Facoetti, 2014), (3) *interactive elements*: incorporating interactive elements to engage learners to explore concepts in a hands-on manner (García-Carrión et al., 2018), and (4) *predictable routine*: establishing a consistent routine for delivering hints consistently can help learners feel more comfortable and confident (Love et al., 2012).

Evaluation Metrics. Evaluation of generated hints is a non-trivial task, depending on multiple factors to determine the quality and the success of a hint. Prior work determines the success of a hint generation system by the learner's abilities in producing the reference solution (Equation (2)) (Jatowt et al., 2023). However, this evaluation framework does not penalize the generated hints if they make the problem too simple, and also does not take into account the individual learner's preferences and learning objectives. Therefore, we need to build human-centered evaluation frameworks (Lee et al., 2022) that can help measure factors beyond the learner's answering capabilities, such as learner's ownership over the learning process, long-term capabilities of generated knowledge, motivation and enjoyment they receive during their interactions.

4.4 Ethics Considerations

The integration of NLP within educational settings raises distinct concerns, such as the impact on pedagogical approaches, the dynamics of teacher-student interaction, and learner agency (Holstein et al., 2019). The adoption of NLP technologies in the classroom implements a particular theory of teaching and learning, and these values must be made explicit (Blodgett and Madaio, 2021). How does introducing a new tool re-configure the dynamics of the teacher-student relationship? Here, it would be crucial to avoid the solutionism trap, define the boundaries of where the system is useful, and ensure that the intention is to augment educators' workflows

instead of substituting them (Remian, 2019). Researchers and practitioners must also attend to the longer-term impacts of engaging with young individuals during a formative period (Holmes et al., 2021). Below, we outline various ethics considerations, including data privacy and consent, bias and fairness, and effects on language variation (Schneider, 2022), and offer strategies to address these concerns.

One of the biggest sets of ethical considerations relates to the use of student and teacher information (Nguyen et al., 2023). Given the sensitive nature of educational data, it will be important to set up privacy measures and enable informed consent. Students' information beyond individual responses to questions may need to be tracked to provide an effective learning experience (Kerr, 2020). However, this opens up the possibility of surveillance and misuse, jeopardizing learners' trust and autonomy (Regan and Steeves, 2019). It would be critical to promote data literacy among educators and learners (e.g., through workshops) to enable them to minimize the risk of their participation (Kerr, 2020). The issue of data ownership raises questions about who holds control over the information collected through education platforms.

We must collectively explore the broader implications of integrating NLP in education on representativeness and equity and exacerbating systemic inequalities (Weidinger et al., 2022). There is a risk that the hints generated by NLP models may not adequately reflect the diverse backgrounds and lived experiences of students (Dixon-Román et al., 2020) and potentially perpetuate harmful stereotypes about different identities (Dev et al., 2021). Prior work has demonstrated the various forms of 'bias' in NLP systems (Blodgett et al., 2020), which may contribute to the construction of language hierarchies and limiting language variation (Schneider, 2022). To promote inclusive design and mitigate these ethical considerations, it is essential to understand how power, privilege, and resources are redistributed as a result of introducing AI in the classroom. Is there a possibility of diminishing quality education for marginalized and under-resourced groups (Remian, 2019)? We must take a community-collaborative approach to understand how to design justice-oriented and accountable systems (Madaio et al., 2022) where learners can truly benefit from hint generation systems.

5 Conclusion

In this paper, we consolidate prior research in hint generation, bridging the gap between research in education and cognitive science, and research in AI and natural language processing. Based on our findings, we propose a roadmap for the future research in hint generation, where we provide a rethinking of the formal task definition, a brief review of research areas that can inform the design of future systems, open challenges for effective hint generation systems, and the ethical considerations. Although hint generation has a long history dating back over three decades (Hume et al., 1996), recent advances in natural language processing could serve useful for future hint generation systems. Beyond education, hint generation is also an excellent atomic task to measure a system's ability to personalize content to user needs and requirements. We invite researchers to foster a community, develop new benchmarks, create shared tasks and workshops for automatic hint generation.

Limitations

Due to the rich and diverse literature on intelligent tutoring systems, we limit our survey to research directly relevant to hint generation and do not cover other types of learning-related feedback.

Acknowledgments

We would like to thank Shivani Kapania for identifying ethics considerations and proofreading this work.

References

M. D. Abdulrahaman, N. Faruk, A. A. Oloyede, N. T. Surajudeen-Bakinde, L. A. Olawoyin, O. V. Mejabi, Y. O. Imam-Fulani, A. O. Fahm, and A. L. Azeez. 2020. Multimedia tools in the teaching and learning processes: A systematic review. *Helijon*, 6(11). <https://doi.org/10.1016/j.helijon.2020.e05312>, PubMed: 33195834

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv: 2303.08774*.

Said Al Faraby, Adiwijaya Adiwijaya, and Ade Romadhony. 2023. Review on neural question generation for education purposes. *International Journal of Artificial Intelligence in Education*. <https://doi.org/10.1007/s40593-023-00374-x>

John R. Anderson. 2013. *The adaptive character of thought*. Psychology Press. <https://doi.org/10.4324/9780203771730>

John R. Anderson, C. Franklin Boyle, and Brian J. Reiser. 1985. Intelligent tutoring systems. *Science*, 228(4698):456–462. <https://doi.org/10.1126/science.228.4698.456>, PubMed: 17746875

John R. Anderson, Albert T. Corbett, Kenneth R. Koedinger, and Ray Pelletier. 1995. Cognitive tutors: Lessons learned. *The Journal of the Learning Sciences*, 4(2):167–207. https://doi.org/10.1207/s15327809jls0402_2

David G. Ausubel. 1963. Cognitive structure and the facilitation of meaningful verbal learning1. *Journal of Teacher Education*, 14(2):217–222. <https://doi.org/10.1177/002248716301400220>

David P. Ausubel. 1962. A subsumption theory of meaningful verbal learning and retention. *The Journal of General Psychology*, 66(2):213–224. <https://doi.org/10.1080/00221309.1962.9711837>, PubMed: 13863333

David Paul Ausubel. 2012. *The acquisition and retention of knowledge: A cognitive view*. Springer Science & Business Media.

Fatemeh Bambaeeroo and Nasrin Shokrpour. 2017. The impact of the teachers' non-verbal communication on success in teaching. *Journal of Advances in Medical Education & Professionalism*, 5(2):51–59. PubMed: 28367460

Tiffany Barnes and John Stamper. 2008. Toward automatic hint generation for logic proof tutoring using historical student data. In *International Conference on Intelligent Tutoring Systems*, pages 373–382. Springer. https://doi.org/10.1007/978-3-540-69132-7_41

Brian R. Belland. 2017. *Instructional Scaffolding in STEM Education: Strategies and Efficacy Evidence*. Springer Nature. <https://doi.org/10.1007/978-3-319-02565-0>

H. Russell Bernard. 1992. Preserving language diversity. *Human Organization*, 51(1):82–89. <https://doi.org/10.17730/humo.51.1.bp4765g377q32761>

Gautam Biswas, James R. Segedy, and John S. Kinnebrew. 2014. –a combined theory-and data-driven approach for interpreting learners’ metacognitive behaviors in open-ended tutoring environments. *Design Recommendations for Intelligent Tutoring Systems*, 2:135.

Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (technology) is power: A critical survey of ‘bias’ in NLP. *arXiv preprint arXiv:2005.14050*.

Su Lin Blodgett and Michael Madaio. 2021. Risks of AI foundation models in education. *arXiv preprint arXiv:2110.10024*.

Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. COMET: Commonsense transformers for automatic knowledge graph construction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4762–4779, Florence, Italy. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1470>

David Boud and Elizabeth Molloy. 2013. Rethinking models of feedback for learning: the challenge of design. *Assessment & Evaluation in Higher Education*, 38(6):698–712. <https://doi.org/10.1080/02602938.2012.691462>

T. M. Bray. 1999. *The Shadow Education System: Private Tutoring and its Implications for Planners*. UNESCO International Institute for Educational Planning.

Monica Bulger. 2016. Personalized learning: The conversations we’re not having. *Data and Society*, 22(1):1–29.

Jannis Bulian, Christian Buck, Wojciech Gajewski, Benjamin Börschinger, and Tal Schuster. 2022. Tomayto, tomahto. Beyond token-level answer equivalence for question answering evaluation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 291–305. <https://doi.org/10.18653/v1/2022.emnlp-main.20>

Deng Cai, Yan Wang, Lemao Liu, and Shuming Shi. 2022. Recent advances in retrieval-augmented text generation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR ’22, pages 3417–3419, New York, NY, USA. Association for Computing Machinery. <https://doi.org/10.1145/3477495.3532682>

Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémie Scheurer, Javier Rando, Rachel Freedman, Tomasz Korbak, David Lindner, Pedro Freire, Tony Wang, Samuel Marks, Charbel-Raphaël Segerie, Micah Carroll, Andi Peng, Phillip Christoffersen, Mehul Damani, Stewart Slocum, Usman Anwar, Anand Siththaranjan, Max Nadeau, Eric J. Michaud, Jacob Pfau, Dmitrii Krasheninnikov, Xin Chen, Lauro Langosco, Peter Hase, Erdem Bıyık, Anca Dragan, David Krueger, Dorsa Sadigh, and Dylan Hadfield-Menell. 2023. Open problems and fundamental limitations of reinforcement learning from human feedback. *arXiv preprint arXiv:2307.15217*.

Steven P. Chamberlain. 2005. Recognizing and responding to cultural differences in the education of culturally and linguistically diverse learners. *Intervention in School and Clinic*, 40(4):195–211. <https://doi.org/10.1177/10534512050400040101>

Kaiyan Chang, Songcheng Xu, Chenglong Wang, Yingfeng Luo, Tong Xiao, and Jingbo Zhu. 2024. Efficient prompting methods for large language models: A survey. *arXiv preprint arXiv:2404.01077*.

Li-Hsin Chang, Jenna Kanerva, and Filip Ginter. 2022. Towards automatic short answer assessment for finnish as a paraphrase retrieval task. In *Proceedings of the 17th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2022)*, pages 262–271. <https://doi.org/10.18653/v1/2022.bea-1.30>

Guanliang Chen, Jie Yang, Claudia Hauff, and Geert-Jan Houben. 2018. Learningq: A large-scale dataset for educational question generation. In *Proceedings of the International*

AAAI Conference on Web and Social Media, volume 12. <https://doi.org/10.1609/icwsm.v12i1.14987>

Michelene T. H. Chi. 1996. Constructing self-explanations and scaffolded explanations in tutoring. *Applied Cognitive Psychology*, 10(7):33–49. [https://doi.org/10.1002/\(SICI\)1099-0720\(199611\)10:7<33::AID-ACP436>3.0.CO;2-E](https://doi.org/10.1002/(SICI)1099-0720(199611)10:7<33::AID-ACP436>3.0.CO;2-E)

Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellar, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.

Janet Collins, Michael Hammond, and Jerry Wellington. 2002. *Teaching and Learning with Multimedia*. Routledge. <https://doi.org/10.4324/9780203441305>

Albert T. Corbett and John R. Anderson. 1994. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-adapted Interaction*, 4:253–278. <https://doi.org/10.1007/BF01099821>

Tyne Crow, Andrew Luxton-Reilly, and Burkhard Wuensche. 2018. Intelligent tutoring systems for programming education: A systematic review. In *Proceedings of the 20th Australasian Computing Education Conference*, pages 53–62. <https://doi.org/10.1145/3160489.3160492>

Marina Danilevsky, Kun Qian, Ranit Aharonov, Yannis Katsis, Ban Kawas, and Prithviraj Sen.

2020. A survey of the state of explainable ai for natural language processing. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 447–459. <https://doi.org/10.18653/v1/2020.aacl-main.46>

Bidyut Das, Mukta Majumder, Santanu Phadikar, and Arif Ahmed Sekh. 2021. Automatic question generation and answer assessment: A survey. *Research and Practice in Technology Enhanced Learning*, 16(1):1–15. <https://doi.org/10.1186/s41039-021-00151-1>

Ana Cláudia Akemi Matsuki de Faria, Felype de Castro Bastos, José Victor Nogueira Alves da Silva, Vitor Lopes Fabris, Valeska de Sousa Uchoa, Décio Gonçalves de Aguiar Neto, and Claudio Filipi Goncalves dos Santos. 2023. Visual question answering: A survey on techniques and common trends in recent literature. *arXiv preprint arXiv:2305.11033*.

Jan Deriu, Alvaro Rodrigo, Arantxa Otegi, Guillermo Echegoyen, Sophie Rosset, Eneko Agirre, and Mark Cieliebak. 2021. Survey on evaluation methods for dialogue systems. *Artificial Intelligence Review*, 54:755–810. <https://doi.org/10.1007/s10462-020-09866-x>, PubMed: 33505103

Sunipa Dev, Emily Sheng, Jieyu Zhao, Aubrie Amstutz, Jiao Sun, Yu Hou, Mattie Sanseverino, Jiin Kim, Akihiro Nishi, Nanyun Peng, and Kai-Wei Chang. 2021. On measures of biases and harms in NLP. *arXiv preprint arXiv:2108.03362*.

Ezekiel Dixon-Román, T. Philip Nichols, and Ama Nyame-Mensah. 2020. The racializing forces of/in ai educational technologies. *Learning, Media and Technology*, 45(3):236–250. <https://doi.org/10.1080/17439884.2020.1667825>

Anmei Dong, Morris Siu-Yung Jong, and Ronnel B. King. 2020. How does prior knowledge influence learning engagement? The mediating roles of cognitive load and help-seeking. *Frontiers in Psychology*, 11:591203. <https://doi.org/10.3389/fpsyg.2020.591203>, PubMed: 33192933

Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey for in-context learning. *arXiv preprint arXiv:2301.00234*.

Jenaro Díaz Ducca. 2014. Positive oral encouragement in the efl classroom, a case study through action research. *Revista de Lenguas Modernas*, 21.

Nicholas Evans and Stephen C. Levinson. 2009. The myth of language universals: Language diversity and its importance for cognitive science. *Behavioral and Brain Sciences*, 32(5):429–448. <https://doi.org/10.1017/S0140525X0999094X>, PubMed: 19857320

Yi Fung, Ruining Zhao, Jae Doo, Chenkai Sun, and Heng Ji. 2024. Massively multi-cultural knowledge acquisition & LM benchmarking. *arXiv preprint arXiv:2402.09369*.

Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*.

Rocío García-Carrión, Silvia Molina Roldán, and Esther Roca Campos. 2018. Interactive learning environments for the educational improvement of students with disabilities in special schools. *Frontiers in Psychology*, 9:1744. <https://doi.org/10.3389/fpsyg.2018.01744>, PubMed: 30283388

Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. Did aristotle use a laptop? A question answering benchmark with implicit reasoning strategies. *Transactions of the Association for Computational Linguistics*, 9:346–361. https://doi.org/10.1162/tacl_a_00370

Carina González, Alberto Mora, and Pedro Toledo. 2014. Gamification in intelligent tutoring systems. In *Proceedings of the Second International Conference on Technological Ecosystems for Enhancing Multiculturality*, pages 221–225. <https://doi.org/10.1145/2669711.2669903>

Simone Gori and Andrea Facoetti. 2014. Perceptual learning as a possible new approach for remediation and prevention of developmental dyslexia. *Vision Research*, 99:78–87. <https://doi.org/10.1016/j.visres.2013.11.011>, PubMed: 24325850

Arthur C. Graesser, Mark W. Conley, and Andrew Olney. 2012. Intelligent tutoring systems. In K. R. Harris, S. Graham, T. Urdan, A. G. Bus, S. Major, & H. L. Swanson (Eds.), *APA educational psychology handbook, Vol. 3. Application to learning and teaching*, pages 451–473. American Psychological Association. <https://doi.org/10.1037/13275-018>

Maureen E. Gray and Keith J. Holyoak. 2021. Teaching by analogy: From theory to practice. *Mind, Brain, and Education*, 15(3):250–263. <https://doi.org/10.1111/mbe.12288>

Alex Guilherme and Ana Lucia Souza de Freitas. 2018. Discussing education by means of metaphors. *Educational Philosophy and Theory*, 50(10):947–956. <https://doi.org/10.1080/00131857.2016.1198250>

Itika Gupta, Barbara Di Eugenio, Devika Salunke, Andrew Boyd, Paula Allen-Meares, Carolyn Dickens, and Olga Garcia. 2020. Heart failure education of African American and Hispanic/Latino patients: Data collection and analysis. In *Proceedings of the First Workshop on Natural Language Processing for Medical Conversations*, pages 41–46, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.nlpmc-1.6>

Telle Hailikari, Nina Katajavuori, and Sari Lindblom-Ylänn. 2008. The relevance of prior knowledge in learning and instructional design. *American Journal of Pharmaceutical Education*, 72(5), 113. <https://doi.org/10.5688/aj7205113>, PubMed: 19214267

Telle Hailikari, Anne Nevgi, and Sari Lindblom-Ylänn. 2007. Exploring alternative ways of assessing prior knowledge, its components and their relation to student achievement: A mathematics based case study. *Studies in Educational Evaluation*, 33(3–4):320–337. <https://doi.org/10.1016/j.stueduc.2007.07.007>

Jennifer Hammond. 2001. *Scaffolding: Teaching and Learning in Language and Literacy Education*. ERIC.

Muhammad Asif Hasan, Nurul Fazmidar Mohd Noor, Siti Soraya Binti Abdul Rahman, and Mohammad Mustaneeer Rahman. 2020. The transition from intelligent to affective tutoring system: A review and open issues. *IEEE Access*, 8:204612–204638. <https://doi.org/10.1109/ACCESS.2020.3036990>

Kilichbek Haydarov, Xiaoqian Shen, Avinash Madasu, Mahmoud Salem, Li-Jia Li, Gamaleldin Elsayed, and Mohamed Elhoseiny. 2023. Affective visual dialog: A large-scale benchmark for emotional reasoning based on visually grounded conversations. *arXiv preprint arXiv:2308.16349*.

Zhicheng He, Weiwen Liu, Wei Guo, Jiarui Qin, Yingxue Zhang, Yaochen Hu, and Ruiming Tang. 2023. A survey on user behavior modeling in recommender systems. *arXiv preprint arXiv:2302.11087*.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8):1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>, PubMed: 9377276

Bas Hofstra, Vivek V. Kulkarni, Sebastian Munoz-Najar Galvez, Bryan He, Dan Jurafsky, and Daniel A. McFarland. 2020. The diversity-innovation paradox in science. *Proceedings of the National Academy of Sciences*, 117(17):9284–9291. <https://doi.org/10.1073/pnas.1915378117>, PubMed: 32291335

Wayne Holmes, Kaska Porayska-Pomsta, Ken Holstein, Emma Sutherland, Toby Baker, Simon Buckingham Shum, Olga C. Santos, Mercedes T. Rodrigo, Mutlu Cukurova, Ig Ibert Bittencourt, and Kenneth R. Koedinger. 2021. Ethics of AI in education: Towards a community-wide framework. *International Journal of Artificial Intelligence in Education*, pages 1–23. <https://doi.org/10.1007/s40593-021-00239-1>

Kenneth Holstein, Bruce M. McLaren, and Vincent Aleven. 2019. Designing for complementarity: Teacher and student needs for orchestration support in ai-enhanced classrooms. In *Artificial Intelligence in Education: 20th International Conference, AIED 2019, Chicago, IL, USA, June 25–29, 2019, Proceedings, Part I* 20, pages 157–171. Springer. https://doi.org/10.1007/978-3-030-23204-7_14

M. Enamul Hoque. 2016. Three domains of learning: Cognitive, affective and psychomotor. *The Journal of EFL Education and Research*, 2(2):45–52.

Jie Huang and Kevin Chen-Chuan Chang. 2022. Towards reasoning in large language models: A survey. *arXiv preprint arXiv:2212.10403*.

Gregory Hume, Joel Michael, Allen Rovick, and Martha Evens. 1996. Hinting as a tactic in one-on-one tutoring. *The Journal of the Learning Sciences*, 5(1):23–47. https://doi.org/10.1207/s15327809jls0501_2

Radhwan Hussein Ibrahim and Dhia-Alrahman Hussein. 2016. Assessment of visual, auditory, and kinesthetic learning style among undergraduate nursing students. *International Journal of Advanced Nursing Studies*, 5(1):1–4. <https://doi.org/10.14419/ijans.v5i1.5124>

Adam Jatowt, Calvin Gehrer, and Michael Färber. 2023. Automatic hint generation. In *Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval*, pages 117–123. <https://doi.org/10.1145/3578337.3605119>

Samantha Jiménez, Reyes Juárez-Ramírez, Victor H. Castillo, and Juan José Tapia Armenta. 2018. *Affective Feedback in Intelligent Tutoring Systems: A Practical Approach*. Springer. <https://doi.org/10.1007/978-3-319-93197-5>

Nengzheng Jin, Joanna Siebert, Dongfang Li, and Qingcai Chen. 2022. A survey on table question answering: Recent advances. In *China Conference on Knowledge Graph and Semantic Computing*, pages 174–186. Springer. https://doi.org/10.1007/978-981-19-7596-7_14

Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. Pub medQA: A dataset for biomedical research question answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural*

Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 2567–2577. <https://doi.org/10.18653/v1/D19-1259>

Wei Jin, Tiffany Barnes, John Stamper, Michael John Eagle, Matthew W. Johnson, and Lorrie Lehmann. 2012. Program representation for automatic hint generation for a data-driven novice programming tutor. In *Intelligent Tutoring Systems: 11th International Conference, ITS 2012, Chania, Crete, Greece, June 14–18, 2012. Proceedings 11*, pages 304–309. Springer. https://doi.org/10.1007/978-3-642-30950-2_40

Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611. <https://doi.org/10.18653/v1/P17-1147>

Azyan Yusra Kapi, Norlis Osman, Ratna Zuarni Ramli, and Jamaliah Mohd Taib. 2017. Multimedia education tools for effective teaching and learning. *Journal of Telecommunication, Electronic and Computer Engineering (JTEC)*, 9(2–8):143–146.

Timo Kaufmann, Paul Weng, Viktor Bengs, and Eyke Hüllermeier. 2023. A survey of reinforcement learning from human feedback. *arXiv preprint arXiv:2312.14925*.

Aniruddha Kembhavi, Minjoon Seo, Dustin Schwenk, Jonghyun Choi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Are you smarter than a sixth grader? Textbook question answering for multimodal machine comprehension. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4999–5007. <https://doi.org/10.1109/CVPR.2017.571>

Kourtney Kerr. 2020. Ethical considerations when using artificial intelligence-based assistive technologies in education. *Ethical Use of Technology in Digital Learning Environments: Graduate Student Perspectives*.

Thomas Kim and Saul Axelrod. 2005. Direct instruction: An educators’ guide and a plea for action. *The Behavior Analyst Today*, 6(2):111. <https://doi.org/10.1037/h0100061>

Donald E. Knuth. 1968. Semantics of context-free languages. *Mathematical Systems Theory*, 2(2):127–145. <https://doi.org/10.1007/BF01692511>

Nizamettin Koc and Bekir Celik. 2015. The impact of number of students per teacher on student achievement. *Procedia-Social and Behavioral Sciences*, 177:65–70. <https://doi.org/10.1016/j.sbspro.2015.02.335>

Martin A. Kozloff, Louis LaNunziata, and James Cowardin. 1999. Direct instruction in education. *Journal Instructivist. Januari*.

Ghader Kurdi, Jared Leo, Bijan Parsia, Uli Sattler, and Salam Al-Emari. 2020. A systematic review of automatic question generation for educational purposes. *International Journal of Artificial Intelligence in Education*, 30:121–204. <https://doi.org/10.1007/s40593-019-00186-y>

Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. Race: Large-scale reading comprehension dataset from examinations. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794. <https://doi.org/10.18653/v1/D17-1082>

Nataša Lalić. 2005. The role of encouragement in primary schools. *Zbornik Instituta za pedagoška istraživanja*, 37(2):132–152. <https://doi.org/10.2298/ZIPI0502132L>

Yunshi Lan, Gaole He, Jinhao Jiang, Jing Jiang, Zhao Wayne Xin, and Ji Rong Wen. 2022. Complex knowledge base question answering: A survey. *IEEE Transactions on Knowledge and Data Engineering*, page 1.

Nguyen-Thinh Le, Sven Strickroth, Sebastian Gross, and Niels Pinkwart. 2013. A review of ai-supported tutoring approaches for learning programming. *Advanced Computational Methods for Knowledge Engineering*, pages 267–279. <https://doi.org/10.1007/978-3-319-00293-4>

Mina Lee, Megha Srivastava, Amelia Hardy, John Thickstun, Esin Durmus, Ashwin Paranjape, Ines Gerard-Ursin, Xiang Lisa Li,

Faisal Ladhak, Frieda Rong, et al. 2022. Evaluating human-language model interaction. *arXiv preprint arXiv:2212.09746*.

Huayang Li, Yixuan Su, Deng Cai, Yan Wang, and Lema Liu. 2022. A survey on retrieval-augmented text generation. *arXiv preprint arXiv:2202.01110*.

Jiazheng Li, Lin Gui, Yuxiang Zhou, David West, Cesare Aloisi, and Yulan He. 2023a. Distilling chatGPT for explainable automated student answer assessment. *arXiv preprint arXiv:2305.12962*.

Jialin Li, Junli Wang, Junjie Hu, and Ming Jiang. 2024. How well do LLMs identify cultural unity in diversity? *arXiv preprint arXiv:2408.05102*.

Qinbin Li, Zeyi Wen, Zhaomin Wu, Sixu Hu, Naibo Wang, Yuan Li, Xu Liu, and Bingsheng He. 2021. A survey on federated learning systems: Vision, hype and reality for data privacy and protection. *IEEE Transactions on Knowledge and Data Engineering*.

Xiaoyu Li. 2021. The application of teachers' encouragement in design classroom. In *Advances in Creativity, Innovation, Entrepreneurship and Communication of Design: Proceedings of the AHFE 2021 Virtual Conferences on Creativity, Innovation and Entrepreneurship, and Human Factors in Communication of Design, July 25–29, 2021, USA*, pages 323–331. Springer. https://doi.org/10.1007/978-3-030-80094-9_39

Yuanchao Li, Koji Inoue, Leimin Tian, Changzeng Fu, Carlos Toshinori Ishi, Hiroshi Ishiguro, Tatsuya Kawahara, and Catherine Lai. 2023b. I know your feelings before you do: Predicting future affective reactions in human-computer dialogue. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–7. <https://doi.org/10.1145/3544549.3585869>

Tzu-Chiang Lin, Ying-Shao Hsu, Shu-Sheng Lin, Maio-Li Changlai, Kun-Yuan Yang, and Ting-Ling Lai. 2012. A review of empirical evidence on scaffolding for science education. *International Journal of Science and Mathematics Education*, 10:437–455. <https://doi.org/10.1007/s10763-011-9322-z>

Chen Liu, Fajri Koto, Timothy Baldwin, and Iryna Gurevych. 2024. Are multilingual LLMs culturally-diverse reasoners? An investigation into multicultural proverbs and sayings. 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 2016–2039. <https://doi.org/10.18653/v1/2024.naacl-long.112>

Haochen Liu, Zitao Liu, Zhongqin Wu, and Jiliang Tang. 2020. Personalized multimodal feedback generation in education. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1826–1840. <https://doi.org/10.18653/v1/2020.coling-main.166>

Naiming Liu, Zichao Wang, Richard Baraniuk, and Andrew Lan. 2022. Open-ended knowledge tracing for computer science education. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3849–3862. <https://doi.org/10.18653/v1/2022.emnlp-main.254>

Jessica J. Love, Caio F. Miguel, Jonathan K. Fernand, and Jillian K. LaBrie. 2012. The effects of matched stimulation and response interruption and redirection on vocal stereotyping. *Journal of Applied Behavior Analysis*, 45(3):549–564. <https://doi.org/10.1901/jaba.2012.45-549>, PubMed: 23060668

Graham Low. 2008. Metaphor and education. *The Cambridge Handbook of Metaphor and Thought*, pages 212–231. <https://doi.org/10.1017/CBO9780511816802.014>

Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521.

Qing Lyu, Marianna Apidianaki, and Chris Callison-Burch. 2024. Towards faithful model explanation in NLP: A Survey. *Computational Linguistics*, pages 1–70. https://doi.org/10.1162/coli_a_00511

Yukun Ma, Khanh Linh Nguyen, Frank Z. Xing, and Erik Cambria. 2020. A survey

on empathetic dialogue systems. *Information Fusion*, 64:50–70. <https://doi.org/10.1016/j.inffus.2020.06.011>

Jakub Macina, Nico Daheim, Sankalan Chowdhury, Tanmay Sinha, Manu Kapur, Iryna Gurevych, and Mrinmaya Sachan. 2023. Mathdial: A dialogue tutoring dataset with rich pedagogical properties grounded in math reasoning problems. *Findings of the Association for Computational Linguistics: EMNLP 2023*. <https://doi.org/10.18653/v1/2023.findings-emnlp.372>

Michael Madaio, Su Lin Blodgett, Elijah Mayfield, and Ezekiel Dixon-Román. 2022. Beyond “fairness”: Structural (in) justice lenses on ai for education. In *The Ethics of Artificial Intelligence in Education*, pages 203–239. Routledge. <https://doi.org/10.4324/9780429329067-11>

Mariam Mahdaoui, N. O. U. H. Said, My Seddiq Elkasmi Alaoui, and Mounir Sadiq. 2022. Comparative study between automatic hint generation approaches in intelligent programming tutors. *Procedia Computer Science*, 198:391–396. <https://doi.org/10.1016/j.procs.2021.12.259>

Jessica McBroom, Irena Koprinska, and Kalina Yacef. 2021. A survey of automated programming hint generation: The hints framework. *ACM Computing Surveys (CSUR)*, 54(8):1–27. <https://doi.org/10.1145/3469885>

John McCarthy. 1964. *A Formal Description of a Subset of ALGOL*. Stanford University. Computer Science Department. Artificial Intelligence. <https://doi.org/10.21236/AD0785050>

Bonan Min, Hayley Ross, Elior Sulem, Amir Pouran Ben Veyseh, Thien Huu Nguyen, Oscar Sainz, Eneko Agirre, Ilana Heintz, and Dan Roth. 2023. Recent advances in natural language processing via large pre-trained language models: A survey. *ACM Computing Surveys*, 56(2):1–40. <https://doi.org/10.1145/3605943>

Michele Miranda, Elena Sofia Ruzzetti, Andrea Santilli, Fabio Massimo Zanzotto, Sébastien Bratières, and Emanuele Rodolà. 2024. Pre-

serving privacy in large language models: A survey on current threats and solutions. *arXiv preprint arXiv:2408.05212*.

Elham Mousavinasab, Nahid Zarifsanaiey, Sharareh R. Niakan Kalhor, Mahnaz Rakhshan, Leila Keikha, and Marjan Ghazi Saeedi. 2021. Intelligent tutoring systems: A systematic review of characteristics, applications, and evaluation methods. *Interactive Learning Environments*, 29(1):142–163. <https://doi.org/10.1080/10494820.2018.1558257>

Jamshid Mozafari, Abdelrahman Abdallah, Bhawna Piriyani, and Adam Jatowt. 2024a. Exploring hint generation approaches in open-domain question answering. *arXiv preprint arXiv:2409.16096*.

Jamshid Mozafari, Anubhav Jangra, and Adam Jatowt. 2024b. Triviahg: A dataset for automatic hint generation from factoid questions. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2060–2070. <https://doi.org/10.1145/3626772.3657855>

Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Naveed Akhtar, Nick Barnes, and Ajmal Mian. 2023. A comprehensive overview of large language models. *arXiv preprint arXiv:2307.06435*.

Andy Nguyen, Ha Ngan Ngo, Yvonne Hong, Belle Dang, and Bich-Phuong Thi Nguyen. 2023. Ethical principles for artificial intelligence in education. *Education and Information Technologies*, 28(4):4221–4241.

Jinjie Ni, Tom Young, Vlad Pandelea, Fuzhao Xue, and Erik Cambria. 2023. Recent advances in deep learning based dialogue systems: A systematic survey. *Artificial Intelligence Review*, 56(4):3055–3155. <https://doi.org/10.1007/s10462-022-10248-8>

Mark Nichter and Mimi Nichter. 2003. Education by appropriate analogy. *Anthropology and International Health*, pages 433–459. Routledge. <https://doi.org/10.4324/9780203393390-25>, <https://doi.org/10.4324/9780203393390>

Florian Obermüller, Ute Heuer, and Gordon Fraser. 2021. Guiding next-step hint generation using automated tests. In *Proceedings of the 26th ACM Conference on Innovation and Technology in Computer Science Education V. 1*, pages 220–226. <https://doi.org/10.1145/3430665.3456344>

Benjamin Paaßen, Barbara Hammer, Thomas Price, Tiffany Barnes, Sebastian Gross, and Niels Pinkwart. 2018. The continuous hint factory—providing hints in vast and sparsely populated edit distance spaces. *Journal of Educational Data Mining*, 10(1).

Sankalan Pal Chowdhury, Vilém Zouhar, and Mrinmaya Sachan. 2024. Autotutor meets large language models: A language model tutor with rich pedagogy and guardrails. In *Proceedings of the Eleventh ACM Conference on Learning@ Scale*, pages 5–15. <https://doi.org/10.1145/3657604.3662041>

Liangming Pan, Wenqiang Lei, Tat-Seng Chua, and Min-Yen Kan. 2019. Recent advances in neural question generation. *arXiv preprint arXiv:1905.08949*.

Baolin Peng, Michel Galley, Pengcheng He, Chris Brockett, Lars Liden, Elnaz Nouri, Zhou Yu, Bill Dolan, and Jianfeng Gao. 2022. GODEL: Large-scale pre-training for goal-directed dialog. *arXiv preprint arXiv:2206.11309*.

Janneke Van de Pol, Monique Volman, and Jos Beishuizen. 2010. Scaffolding in teacher-student interaction: A decade of research. *Educational Psychology Review*, 22:271–296. <https://doi.org/10.1007/s10648-010-9127-6>

Arina Pourkamali, Akbar Mohammadi, and Sara Haghigiat. 2021. The effect of education based on fernald’s multisensory approach on improving visual memory and fluency of students with learning disabilities. *Journal of Adolescent and Youth Psychological Studies (JAYPS)*, 2(2):290–298. <https://doi.org/10.61838/kman.jayps.2.2.24>

Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah A. Smith, and Mike Lewis. 2022. Measuring and narrowing the compositionality gap in language models. *arXiv preprint arXiv:2210.03350*.

Thomas W. Price, Yihuan Dong, and Tiffany Barnes. 2016. Generating data-driven hints for open-ended programming. *International Educational Data Mining Society*.

Thomas W. Price, Yihuan Dong, and Dragan Lipovac. 2017. iSNAP: Towards intelligent tutoring in novice programming environments. In *Proceedings of the 2017 ACM SIGCSE Technical Symposium on Computer Science Education*, pages 483–488. <https://doi.org/10.1145/3017680.3017762>

Thomas W. Price, Yihuan Dong, Rui Zhi, Benjamin Paaßen, Nicholas Lytle, Veronica Cateté, and Tiffany Barnes. 2019. A comparison of the quality of data-driven programming hint generation algorithms. *International Journal of Artificial Intelligence in Education*, 29:368–395. <https://doi.org/10.1007/s40593-019-00177-z>

Jingyuan Qi, Zhiyang Xu, Ying Shen, Minqian Liu, Di Jin, Qifan Wang, and Lifu Huang. 2023. The art of socratic questioning: Recursive thinking with large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4177–4199. <https://doi.org/10.18653/v1/2023.emnlp-main.255>

Aravind Sesagiri Raamkumar and Yinping Yang. 2022. Empathetic conversational systems: A review of current advances, gaps, and opportunities. *IEEE Transactions on Affective Computing*, 14(4):2722–2739. <https://doi.org/10.1109/TAFFC.2022.3226693>

Mirjana Radovic-Markovic and Dusan Markovic. 2012. A new model of education: Development of individuality through the freedom of learning. *Eruditio, e-Journal of the World Academy of Art & Science*, 1(1).

Arief Ramadhan, Harco Leslie Hendric Spits Warnars, and Fariza Hanis Abdul Razak. 2024. Combining intelligent tutoring systems and gamification: A systematic literature review. *Education and Information Technologies*, 29(6):6753–6789. <https://doi.org/10.1007/s10639-023-12092-x>

Priscilla Regan and Valerie Steeves. 2019. Education, privacy and big data algorithms: Taking the persons out of personalized learning. *First Monday*. <https://doi.org/10.5210/fm.v24i11.10094>

Dana Remian. 2019. Augmenting education: Ethical considerations for incorporating artificial intelligence in education.

Lindsey Engle Richland and Nina Simms. 2015. Analogy, higher order thinking, and education. *Wiley Interdisciplinary Reviews: Cognitive Science*, 6(2):177–192. <https://doi.org/10.1002/wcs.1336>, PubMed: 26263071

Peter Rillero. 2016. Deep conceptual learning in science and mathematics: Perspectives of teachers and administrators. *The Electronic Journal for Research in Science & Mathematics Education*, 20(2).

Kelly Rivers and Kenneth R. Koedinger. 2014. Automating hint generation with solution space path construction. In *Intelligent Tutoring Systems: 12th International Conference, ITS 2014, Honolulu, HI, USA, June 5–9, 2014. Proceedings 12*, pages 329–339. Springer. https://doi.org/10.1007/978-3-319-07221-0_41

Kelly Rivers and Kenneth R. Koedinger. 2017. Data-driven hint generation in vast solution spaces: A self-improving python programming tutor. *International Journal of Artificial Intelligence in Education*, 27:37–64. <https://doi.org/10.1007/s40593-015-0070-z>

Reudismam Rolim, Gustavo Soares, Loris D’Antoni, Oleksandr Polozov, Sumit Gulwani, Rohit Gheyi, Ryo Suzuki, and Björn Hartmann. 2017. Learning syntactic program transformations from examples. In *2017 IEEE/ACM 39th International Conference on Software Engineering (ICSE)*, pages 404–415. IEEE. <https://doi.org/10.1109/ICSE.2017.44>

Barak Rosenshine. 2008. Five meanings of direct instruction. *Center on Innovation & Improvement, Lincoln*, pages 1–10.

Pranab Sahoo, Ayush Kumar Singh, Sriparna Saha, Vinija Jain, Samrat Mondal, and Aman Chadha. 2024. A systematic survey of prompt engineering in large language models: Techniques and applications. *arXiv preprint arXiv:2402.07927*.

Britta Schneider. 2022. Multilingualism and AI: The regimentation of language in the age of digital capitalism. *Signs and Society*, 10(3):362–387. <https://doi.org/10.1086/721757>

Gesina Schwalbe and Bettina Finzel. 2023. A comprehensive taxonomy for explainable artificial intelligence: A systematic survey of surveys on methods and concepts. *Data Mining and Knowledge Discovery*, pages 1–59. <https://doi.org/10.1007/s10618-022-00867-8>

Anna Sfard. 2012. Metaphors in education. In *Educational Theories, Cultures and Learning*, pages 39–49. Routledge. <https://doi.org/10.4324/9780203379417-10>, PubMed: 21668129

Veysel Sönmez. 2017. Association of cognitive, affective, psychomotor and intuitive domains in education, sönmez model. *Universal Journal of Educational Research*, 5(3):347–356. <https://doi.org/10.13189/ujer.2017.050307>

Lourdes Diaz Soto, Jocelynn L. Smrekar, and Deanna L. Nekcovei. 1999. Preserving home languages and cultures in the classroom: Challenges and opportunities. *Directions in Language and Education*.

Yash Srivastava, Vaishnav Murali, Shiv Ram Dubey, and Snehasis Mukherjee. 2021. Visual question answering using deep learning: A survey and performance analysis. In *Computer Vision and Image Processing: 5th International Conference, CVIP 2020, Prayagraj, India, December 4–6, 2020, Revised Selected Papers, Part II 5*, pages 75–86. Springer. https://doi.org/10.1007/978-981-16-1092-9_7

Sanja Štajner. 2021. Automatic text simplification for social good: Progress and challenges. *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2637–2652. <https://doi.org/10.18653/v1/2021.findings-acl.233>

Rahmah Wahdaniati Suaib. 2017. The use of visual auditory kinesthetic (vak) learning styles to increase students’ vocabulary. *Didaktika Jurnal Kependidikan*, 11(2):239–253. <https://doi.org/10.30863/didaktika.v11i2.169>

Osama Swidan and Elena Naftaliev. 2019. The role of the design of interactive diagrams in teaching–learning the indefinite

integral concept. *International Journal of Mathematical Education in Science and Technology*, 50(3):464–485. <https://doi.org/10.1080/0020739X.2018.1522674>

Anaïs Tack and Chris Piech. 2022. The AI teacher test: Measuring the pedagogical ability of blender and gpt-3 in educational dialogues. *arXiv preprint arXiv:2205.07540*.

Paul Thagard. 1992. Analogy, explanation, and education. *Journal of Research in Science Teaching*, 29(6):537–544. <https://doi.org/10.1002/tea.3660290603>

Christine D. Tippett. 2016. What recent research on diagrams suggests about learning *with* rather than learning *from* visual representations in science. *International Journal of Science Education*, 38(5):725–746. <https://doi.org/10.1080/09500693.2016.1158435>

George Tsatsaronis, Georgios Balikas, Prodromos Malakasiotis, Ioannis Partalas, Matthias Zschunke, Michael R. Alvers, Dirk Weissenborn, Anastasia Krithara, Sergios Petridis, Dimitris Polychronopoulos, Yannis Almirantis, John Pavlopoulos, Nicolas Baskiotis, Patrick Gallinari, Thierry Artiéres, Axel-Cyrille Ngonga Ngomo, Norman Heino, Eric Gaussier, Liliana Barrio-Alvers, Michael Schroeder, Ion Androutsopoulos, and Georgios Paliouras. 2015. An overview of the bioasq large-scale biomedical semantic indexing and question answering competition. *BMC Bioinformatics*, 16(1):1–28. <https://doi.org/10.1186/s12859-015-0564-6>, PubMed: 25925131

Yvonne Vezzoli, Asimina Vasalou, and Kaska Porayska-Pomsta. 2017. Dyslexia in sns: An exploratory study to investigate expressions of identity and multimodal literacies. *Proceedings of the ACM on Human-computer Interaction*, 1(CSCW):1–14. <https://doi.org/10.1145/3134738>

L. S. Vygotsky. 1978. *Mind in Society: Development of Higher Psychological Processes*. Harvard University Press.

Akhtim Wahyuni. 2018. The power of verbal and nonverbal communication in learning. In *1st International Conference on Intellectuals' Global Responsibility (ICIGR 2017)*, pages 80–83. Atlantis Press. <https://doi.org/10.2991/icigr-17.2018.19>

Lucy Lu Wang, Kyle Lo, Yoganand Chandrasekhar, Russell Reas, Jiangjiang Yang, Douglas Burdick, Darrin Eide, Kathryn Funk, Yannis Katsis, Rodney Kinney, Yunyao Li, Ziyang Liu, William Merrill, Paul Mooney, Dewey Murdick, Devvret Rishi, Jerry Sheehan, Zhihong Shen, Brandon Stilson, Alex Wade, Kuansan Wang, Nancy Xin Ru Wang, Chris Wilhelm, Boya Xie, Douglas Raymond, Daniel S. Weld, Oren Etzioni, and Sebastian Kohlmeier. 2020. Cord-19: The covid-19 open research dataset. *ArXiv preprint arXiv:2004.10706*.

Rose E. Wang, Qingyang Zhang, Carly Robinson, Susanna Loeb, and Dorottya Demszky. 2023. Step-by-step remediation of students' mathematical mistakes. *arXiv preprint arXiv:2310.10648*.

Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.

Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, Courtney Biles, Sasha Brown, Zac Kenton, Will Hawkins, Tom Stepleton, Abeba Birhane, Lisa Anne Hendricks, Laura Rimell, William Isaac, Julia Haas, Sean Legassick, Geoffrey Irving, and Iason Gabriel. 2022. Taxonomy of risks posed by language models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 214–229. <https://doi.org/10.1145/3531146.3533088>

Johannes Welbl, Nelson F. Liu, and Matt Gardner. 2017. Crowdsourcing multiple choice science questions. *arXiv preprint arXiv:1707.06209*.

William Winn. 1991. Learning from maps and diagrams. *Educational Psychology Review*, 3:211–247. <https://doi.org/10.1007/BF01320077>

David Wood, Jerome S. Bruner, and Gail Ross. 1976. The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry*, 17(2):89–100. <https://doi.org/10.1111/j.1469-7610.1976.tb00381.x>, PubMed: 932126

Mengmeng Yang, Taolin Guo, Tianqing Zhu, Ivan Tjuawinata, Jun Zhao, and Kwok-Yan Lam. 2023. Local differential privacy and its applications: A comprehensive survey. *Computer Standards & Interfaces*, page 103827. <https://doi.org/10.1016/j.csi.2023.103827>

Yu-Hsi Yuan, Ming-Hsiung Wu, Meng-Lei Hu, and I-Chien Lin. 2019. Teacher’s encouragement on creativity, intrinsic motivation, and creativity: The mediating role of creative process engagement. *The Journal of Creative Behavior*, 53(3):312–324. <https://doi.org/10.1002/jocb.181>

Michael Zhang and Eunsol Choi. 2021. SituatedQA: Incorporating extra-linguistic contexts into QA. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7371–7387. <https://doi.org/10.18653/v1/2021.emnlp-main.586>

Ruqing Zhang, Jiafeng Guo, Lu Chen, Yixing Fan, and Xueqi Cheng. 2021. A review on question generation from natural language text. *ACM Transactions on Information Systems (TOIS)*, 40(1):1–43. <https://doi.org/10.1145/3468889>

Yiqun Zhang, Xiaocui Yang, Xingle Xu, Zeran Gao, Yijie Huang, Shiyi Mu, Shi Feng, Daling Wang, Yifei Zhang, Kaisong Song, and Ge Yu. 2024. Affective computing in the era of large language models: A survey from the NLP perspective. *arXiv preprint arXiv:2408.04638*.

Zhuosheng Zhang, Yao Yao, Aston Zhang, Xiangru Tang, Xinbei Ma, Zhiwei He, Yiming Wang, Mark Gerstein, Rui Wang, Gongshen Liu, and Hai Zhao. 2023. Igniting language intelligence: The hitchhiker’s guide from chain-of-thought reasoning to language agents. *arXiv preprint arXiv:2311.11797*.

Zhuosheng Zhang and Hai Zhao. 2021. Advances in multi-turn dialogue comprehension: A survey. *arXiv preprint arXiv:2103.03125*.

Ying Zhao and Jinjun Chen. 2022. A survey on differential privacy for unstructured data content. *ACM Computing Surveys (CSUR)*, 54(10s):1–28. <https://doi.org/10.1145/3490237>

Yaoyao Zhong, Wei Ji, Junbin Xiao, Yicong Li, Weihong Deng, and Tat-Seng Chua. 2022. Video question answering: Datasets, algorithms and challenges. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6439–6455. <https://doi.org/10.18653/v1/2022.emnlp-main.432>

Kurtis Zimmerman and Chandan R. Rupakheti. 2015. An automated framework for recommending program elements to novices (n). In *2015 30th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 283–288. IEEE. <https://doi.org/10.1109/ASE.2015.54>

Alex Zurek, Julia Torquati, and Ibrahim Acar. 2014. Scaffolding as a tool for environmental education in early childhood. *International Journal of Early Childhood Environmental Education*, 2(1):27–57.