

I2EDL: Interactive Instruction Error Detection and Localization

Francesco Taioli^{1,4}, Stefano Rosa², Alberto Castellini¹, Lorenzo Natale²,
Alessio Del Bue², Alessandro Farinelli¹, Marco Cristani¹, Yiming Wang³

"Exit through the door in front of the pool table. Turn left and go behind the sofa. Then, go straight through the door between the table and the windows. Go straight until you reach a circular table, then turn *left*. Go straight and stop in front of the *lamp*"

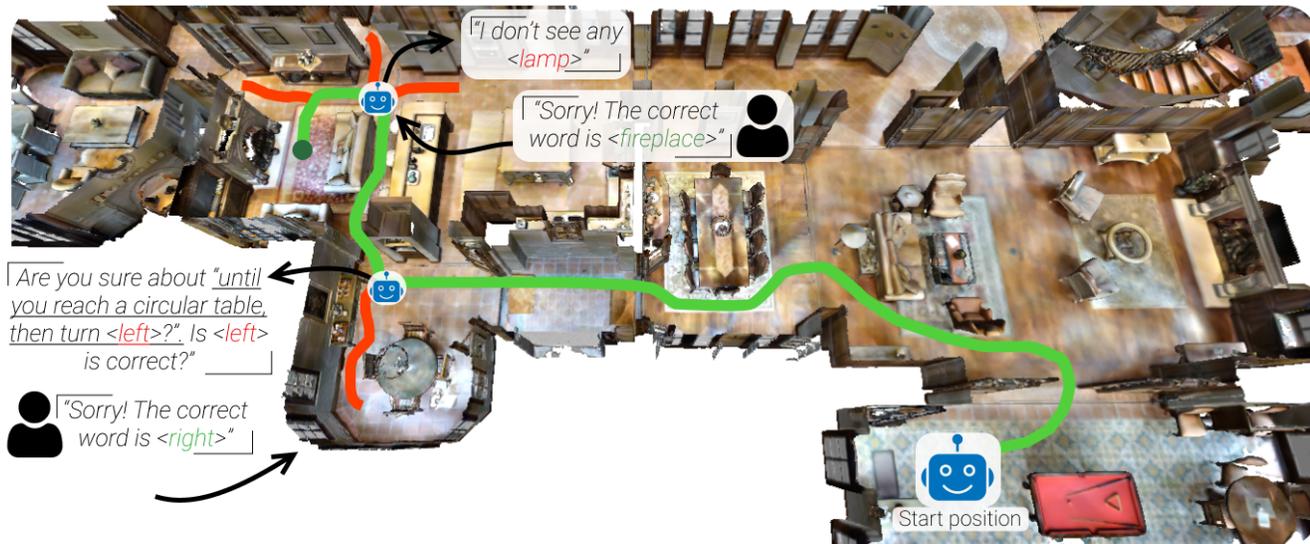


Fig. 1: A human user guides an agent (bottom right) to reach a target goal (green path) with instructions expressed in natural language (top) and which may contain errors (red words). During the navigation, the agent is able to detect and localize instruction errors. Upon detection, the agent asks the user if a particular word in the instruction is correct or not. In case of an incorrect word, the user can reply with the correct one, allowing the agent to resume the navigation. Such human-agent interaction can occur multiple times depending on the error detection algorithm.

Abstract—In the Vision-and-Language Navigation in Continuous Environments (VLN-CE) task, the human user guides an autonomous agent to reach a target goal via a series of low-level actions following a textual instruction in natural language. However, most existing methods do not address the likely case where users may make mistakes when providing such instruction (e.g., “turn left” instead of “turn right”). In this work, we address a novel task of *Interactive VLN in Continuous Environments (IVLN-CE)*, which allows the agent to interact with the user during the VLN-CE navigation to verify any doubts regarding the instruction errors. We propose an *Interactive*

Instruction Error Detector and Localizer (I2EDL) that triggers the user-agent interaction upon the detection of instruction errors during the navigation. We leverage a pre-trained module to detect instruction errors and pinpoint them in the instruction by cross-referencing the textual input and past observations. In such way, the agent is able to query the user for a timely correction, without demanding the user’s cognitive load, as we locate the probable errors to a precise part of the instruction. We evaluate the proposed I2EDL on a dataset of instructions containing errors, and further devise a novel metric, the *Success weighted by Interaction Number (SIN)*, to reflect both the navigation performance and the interaction effectiveness. We show how the proposed method can ask focused requests for corrections to the user, which in turn increases the navigation success, while minimizing the interactions.

I. INTRODUCTION

The research of Vision-and-Language Navigation (VLN) aims to develop agents that can navigate to a specified location within a 3D space by following instructions expressed in natural language. This research task aligns with the broader ambition of embodied AI, which allows automated agents to

¹ University of Verona, Verona, Italy.

² Istituto Italiano di Tecnologia (IIT), Genova, Italy.

³ Fondazione Bruno Kessler, Trento, Italy.

⁴ Polytechnic of Turin, Turin, Italy.

francesco.taioli@polito.it

We acknowledge the CINECA award under the IS CRA initiative, for the availability of high-performance computing resources and support. This work was partially sponsored by the PNR project FAIR - Future AI Research (PE00000013), under the NRRP MUR program funded by the NextGenerationEU. This project has received funding from the European Union’s Horizon research and innovation programme G.A. n. 101070227 (CONVINCE).

engage with human users via natural language conversations, understanding their surroundings [1], and executing tasks in the real world. There are many benchmark datasets in the literature of VLN, including the seminal dataset Room-to-Room (R2R) for discrete environments operating on discrete navigation graphs [2], and the more realistic R2R for continuous environments (R2R-CE) operating via a set of low-level actions to any unobstructed point in a scene [3]. Most of the benchmark datasets consider only *correct* language instructions, implying that the users never make mistakes. However, this is not always the case, since instructions can be very complex and the ability of giving right directions vary greatly among people [4]. In fact, it is a matter of *spatial cognition* skills, i.e., how humans mentally represent and process spatial information [5], and the ability of creating *cognitive maps*, which are mental representation of the layout and contents of an environment [6].

Recently, Taioli *et al.* propose the R2RIE-CE benchmark dataset [7], which introduces wrong instructions in the form of incorrect directions (left, right, etc.) and misplaced rooms or objects, to model the effect of inaccurate cognitive maps. This is helpful in benchmarking the robustness of state-of-the-art VLN policies, that have been designed to cope with ideal interactions only. The R2RIE-CE dataset can also be used to evaluate algorithms for detecting and localizing errors in a given instruction. Yet, their detection and localization baselines operate in an offline mode, i.e., the errors in the instructions are individuated only after an agent has finished its search, thus leaving no chance for the agent and the user to interact and recover the errors *while exploring the environment*. On the contrary, enabling human-agent interaction during navigation could be effective. As shown in Fig. 1, a human user may initially give an erroneous instruction, with a wrong direction word “*left*” to “*right*” and a wrong object landmark “*lamp*” to “*fireplace*”, which can cause the agent to deviate from its target position, failing the navigation task. However, if the agent can detect and locate potential errors while navigating and observing the scene, it can prompt the user for instruction corrections, thus improving the success rate of the task. Such interactive VLN with error awareness introduces additional challenges on top of the existing VLN task (e.g., visual perception, spatial reasoning and vision-language alignment). In particular, the agent should identify potential errors promptly at an early stage, with only partial observations of the scene. Moreover, since it is not ideal to have an agent that constantly interacts with a human user asking for potential errors (both for human disturbance and cognitive load), it is essential to have an accurate online instruction error detector and localizer, thus “*asking the right question at the right time.*”

In this work, we address *Interactive VLN in Continuous Environments (IVLN-CE)*. Human users are allowed to make errors in their initial instructions and subsequently correct them if the agent accurately detects and locates the errors through human-agent interactions. We propose an effective baseline, named Interactive Instruction Error Detector and Localizer (*I2EDL*), that operates in an online mode given

only partial observations. Inspired by the offline method IEDL proposed in [7], we first collect a set of visual observations from the agent. By leveraging the pre-trained models in IEDL, we can identify errors within the instruction and precisely locate them. Then, upon positive detection, the agent asks the user if a specific word is wrong and, if this is the case, it obtains an accurate replacement. In addition to common VLN performance metrics, i.e., Success Rate (SR) and Success weighted by Path Length (SPL), we also propose a novel metric that is specific for the *IVLN-CE* task, the *Success weighted by Interaction Number (SIN)*, which reflects both the navigation performance and interaction effectiveness, by encouraging a higher success rate while limiting the interaction numbers. We evaluate our method on R2RIE-CE under different instruction errors and prove that our baseline is more effective than an agent that randomly interacts with the user. In summary, our contributions are listed below:

- We establish the *IVLN-CE* task, i.e., *Interactive VLN in Continuous Environments*, simulating the real world cases where humans are allowed to make mistakes when providing instruction, and agents are allowed to interact with humans to correct them.
- We propose an effective baseline, *I2EDL*, which interacts with the user in an online manner upon detecting instruction errors and prompting the focused question with localized errors.
- We propose a novel metric that measure the interaction effectiveness in terms of navigation performance by combining the Success Rate and interaction numbers with the user. Our metric serves as a primary quantifier for comparing the performance of different agents.

II. RELATED WORKS

Vision and Language Navigation. The task of *Vision-and-Language Navigation (VLN)* was initially introduced in [2]. Early iterations were based on the Matterport3D Sim [8] and on the Room-to-Room (R2R) dataset [2], which represents the environment as a sparse undirected graph of poses with associated panoramic views. In this discrete environment, agents can only move between pre-existing nodes of the graph. A continuous environment variant of the task (VLN-CE), together with a new dataset (R2R-CE) was introduced in [3], using the Habitat simulator [9]. In VLN-CE, agents are required to follow the instructions while navigating freely in the environment rather than teleporting between nodes. While initial approaches used recurrent representations to encode the agent’s history and predict the next action [10], more recent methods proposed to predict candidate waypoints and select the next best goal [11], the use of topological memory [12] to better encode the history of observations, and metric map pre-training [13] to increase spatial awareness. Recent works have studied how VLN agents use directions and objects to navigate by masking words [14], [15] as well as the performance drop in unseen environments [16]. However, current methods and datasets do not consider the case of

potentially wrong instructions. We propose to specifically study the case of instructions containing errors.

Cooperative Vision-dialog Navigation. The task of *Navigation from Dialogue History* (NDH) was introduced in [17]. In NDH, an agent is given a target object and a dialogue history between two humans cooperating to find it. These approaches are usually evaluated only in terms of progress towards the goal. [18] propose turn-based dialogue between two agents: the navigating agent and the guiding agent. Both agents learn to simulate questions or answers by the other. Vision-dialog navigation has been extended to the real world. [19] proposed RobotSlang, a dataset of natural language dialogues between a human operator teleoperating a robot and a human commander that provides guidance towards a goal. In [20] the agent is trained to identify when to engage in dialogue with a navigator agent via masking and directional grounding. [21] proposes to factor out the action of querying into two different whether-to-ask and a what-to-ask policies.

Interactive VLN. More closely related to our proposal are approaches where the agent interacts with an oracle by asking for help. In [22] the agent uses a dedicated policy action for asking for help. The action of querying the oracle is based on model confusion (*i.e.*, the agent is unsure about which action to take next) and is penalized via a negative reward. When queried, the oracle returns the next shortest path action to the goal. A metric is also introduced to evaluate the effectiveness of human-agent interaction, as the percentage of total ask actions per episode. [22] injects a probability of the oracle making a mistake, to simulate a more realistic user. We directly start with a dataset containing mistakes, to simulate wrong instructions given to the agent. A different way to limit interactions with the oracle is fixing a query budget. In [23] the agent asks for help when unsure about the next action or lost. Upon being called, the oracle provides a short-term goal in natural language. Intervention could be direct (the oracle takes control of the agent) or indirect (the oracle adds new information via short-term textual instructions). A dedicated policy is trained to ask for intervention based on the budget. [24] address the task of *Audio-Visual-Language Embodied Navigation* (AVLEN). The agent can query an oracle under a budget (*i.e.*, the maximum number of queries is limited). The effect of number of interactions on success is not directly evaluated, but only indirectly through success metrics. [25] relaxes assumptions on the oracle by simulating assistants that are only aware of the agent (and can thus provide assistance) when it enters their zone of attention. Interaction is evaluated in terms of number of requests per task. In both [24] and [25], the oracle replies in the form of a full textual instruction. In contrast, in our proposal, the oracle only substitutes a wrong word with the correct one in the original instruction.

III. OUR CONTRIBUTIONS

Task Formulation of IVLN-CE. For each episode i , a human describes to an agent how to reach a target goal by means of a natural language instruction \mathcal{I}_i composed of F words, *i.e.*, $\mathcal{I}_i = \{w_1, \dots, w_F\}$. Note that the instruction given by the human may contain errors, formally, instruction

\mathcal{I}_i can have up to E words that are incorrect. We then define the instruction embedding as $\Upsilon_i \in \mathbb{R}^{W \times D}$, *i.e.*, the instruction \mathcal{I}_i is tokenized and padded up to $W = 80$ tokens, while D is the dimension of the latent space in which each token is projected, following [12], [13]. Without loss of generality, we assume that each word is tokenized into one token.

At each time step t , the agent receives a visual observation O_t , namely an RGB-D image. Let T be the total number of steps executed by policy π , we define the set of visual observation $\mathcal{O} = \{O_1, \dots, O_T\}$. Policy π , for every step t , predicts an action a_t in the set $\{\text{Forward } 0.25\text{m}, \text{Turn Left } 15^\circ, \text{Turn Right } 15^\circ, \text{Stop}\}$.

For every episode i , and at every step t , the agent has the possibility to query the human, checking if a particular token ℓ_i^j is correct or not, where $j \in [0, \text{len}(\Upsilon_i) - 1]$ and $\text{len}(\cdot)$ returns the total number of tokens for instruction Υ_i .

Asking just a token to the user would be ineffective, since the user would hardly understand the sense of a single word (token), and, in the case of multiple instances of the same word, misunderstandings could easily arise. Therefore, the agent passes to the users a portion (*context*) of the instruction, made by multiple tokens $L_{j, \varsigma_l} = [\ell_i^{j-\varsigma_l}, \ell_i^{j+\varsigma_l}]$, where ς_l is the *contextualization length*. The human, upon receiving a request from the agent, returns the real token if a wrong token is found within $[\ell_i^{j-\tau_l}, \ell_i^{j+\tau_l}]$, where τ_l is a *localization threshold*. This correction mechanism ensures that the human can provide the correct token even if there is a slight discrepancy of τ_l tokens in the location pointed out by the agent.

Policy. The agent’s policy π is implemented in this paper by the current state-of-the-art method for VLN-CE¹, *i.e.*, BEVBert [13]. However, our approach is model-agnostic. BEVBert is composed of three essential modules that allow the agent to balance the demand for short-term reasoning and long-term planning: (i) a graph-based topological map for long-term planning equipped with a global action space; (ii) a local metric map for short-term reasoning equipped with a local action space; (iii) a map-based pre-training paradigm to enhance the spatial-aware cross-modal reasoning. Formally, given an episode i , let t be the current time step and n the current node of the topological graph inside the environment. Policy π selects the best candidate node from the topological map of point (i), and low-level actions are performed to bring the agent from node n_t to n_{t+1} . To be comparable to other VLN-CE agents, we maintain a maximum number of $k = 15$ steps, as done in BEVBert [13]. Notice that we do not train the policy π , as it is considered as given. The focus of our work is user-agent interactions during navigation when any instruction error is detected and localized in an online manner.

Instruction Error Detection & Localization. We employ the recently proposed Instruction Error Detector & Localizer (IEDL) [7]. IEDL is firstly composed of a cross-modal transformer, which fuses together the semantic meaning of the language instruction with the sequence of the visual observations of the agent, producing visual-language-aligned features. Then, these features are fed to two heads: (i) a

¹Eval AI - VLN-CE Challenge

detection head f_d , trained to detect when the instruction does not align with the sequence of observations, which outputs an alignment score $a \in [0, 1]$; (ii) a localization head f_l , which predicts the locations of the words that may introduce errors within the instructions. Notably, *IEDL* is trained and evaluated in an offline manner where errors are detected and localized after each complete episode.

I2EDL: Interactive-IEDL. For each episode i , we execute the policy π following instruction \mathcal{I}_i for at least p steps to acquire the set of visual observations $\mathcal{O} = \{O_1, \dots, O_p\}$. When the current step $t \geq p$, we use the detection head f_d of *IEDL* to check if the alignment score is $a \geq \tau_d$, where $\tau_d = 0.6$ is a *detection threshold*. If the detection is positive, meaning that the instruction contains at least one error, we use the *IEDL* localization head f_l to localize the errors. Formally, we apply the *softmax* operator over the output of f_l , and then select the token ℓ_i^j with the highest probability, where j is the index of the token. When a positive detection occurs, we increment variable NI_i , showing that the agent has detected and localized errors in the instruction. We then simulate an agent-human interaction by asking the question “*I think there is an error in this part of the instruction: <part>, and specifically on this <token>. Is this the case?*” In this question, $\langle token \rangle$ refers to token ℓ_i^j , while $\langle part \rangle$ refers to context L_{j, s_l} . If the range identified by tokens $[\ell_i^{j-\tau_l}, \ell_i^{j+\tau_l}]$ contains the error, then the agent receives the correct token, and the embedding for instruction Υ_i is recomputed. Otherwise, if the detection is a false positive or no error is found by the human within the specified range, no action is performed by the human. After the interaction, the agent resumes its navigation, having the possibility to query the user for every step until T steps are performed or action `STOP` action is selected by the policy, meaning that the agent believes to have reached the goal.

Success weighted by Interaction Number. Since our interaction scheme for VLN is new, we need to propose a novel figure of merit. The rationale is that we want to weight the success rate, dependently on how many times the agent requires the human intervention: the higher the number of interventions, the less valuable the success rate. We thus propose *SIN*, *i.e.*, *Success weighted by Interaction Number*, specifically designed to combine, in a single measure, both the *SR* and the number of interactions with the user. Inspired by the *Success weighted by Path Length* metric [26], we define *SIN* as:

$$SIN = \frac{1}{N} \sum_{i=1}^N S_i \frac{1}{1 + \lambda \frac{NI_i}{\max(1, NE_i)}} \quad (1)$$

where NI_i is the number of interactions with the user, NE_i is the number of errors in episode i and S_i is a binary indicator of success for episode i . The $\max(\cdot)$ operator at the denominator ensures the number of interactions NI_i are weighted by the number of errors NE_i . Note that if no errors are present for episode i , $\max(1, NE_i) = 1$. λ is a weighting factor that modulates the penalisation for the interaction numbers. We consider *SIN* as the primary metric in evaluating methods addressing the *IVLN-CE* task.

SIN properties. *SIN* is ranged between 0 and 1. A higher value indicates a better navigation performance with interaction efficiency. Moreover, the proposed *SIN* metric possesses several favourable properties:

(i) when no interaction is performed with the human (*i.e.*, when NI_i is 0), *SIN* is mathematically equivalent to *SR*.

(ii) *SIN* penalizes false positives detections.

Proof: for every correct episode i (*i.e.*, the instruction is correct and thus $NE_i = 0$), a perfect agent will not interact with the human, thus $NI_i = 0$. If this is not the case, NI_i is increased accordingly, thus minimizing the *SIN* metric. Note that, in this scenario, the denominator is 1, resulting in increased importance assigned to each unnecessary interaction.

(iii) *SIN* penalizes repetitive interactions.

Proof: for every incorrect episode i (*i.e.*, the instructions contains error), the *SIN* metric will be penalized as the agent request multiple interactions with the human.

(iv) the weighting factor λ prevents the denominator from becoming excessively large. The metric can still show the improvement in *SR* while penalizing excessive interaction. We found that a $\lambda = 0.01$ is a good compromise between weighting *SR* and number of interactions.

IV. EXPERIMENTS

Dataset. We evaluate our method *I2EDL* in the recently proposed *R2RIE-CE* dataset [7], *i.e.*, *R2R* with *Instruction Errors in Continuous Environment*. *R2RIE-CE* is composed of five sets by incorporating various types of instruction errors, including: (i) *Direction* (one error), (ii) *Object* (one error), (iii) *Room* (one error), (iv) *Room&Object* (two errors); and finally (v) *All* (three errors). Each set \mathcal{E} is composed of correct episodes \mathcal{E}_c and perturbed episodes \mathcal{E}_p . For each episode $e_i \in \mathcal{E}_c$, the authors derived an associated perturbed episode containing specific instruction errors, which is stored in the associated set \mathcal{E}_p . The ratio of \mathcal{E}_c and \mathcal{E}_p is 50%.

Metrics. To appreciate the qualities of our proposed metric *SIN* *w.r.t.* other, standard, VLN metrics [2], [26], we consider: (i) *Success Rate (SR)*: an episode is considered successful if the distance between the final position of the agent and the target location is less than 3 meters. (ii) *Success weighted by Path Length* [26] (*SPL*): defined as

$$SPL = \frac{1}{N} \sum_{i=1}^N S_i \frac{\ell_i}{\max(p_i, \ell_i)}$$

where N are the total number of episodes \mathcal{E} , ℓ_i is the shortest path distance from the agent’s starting position to the goal in episode i , and p_i is the length of the path actually taken by the agent in episode i and S_i is the binary indicator of success for episode i . Finally, (iii) we report *Mean Interaction Number (MIN)*, which is defined as

$$MIN = \frac{1}{N} \sum_{i=1}^N NI_i$$

where NI_i is the number of times the agent interacts with the user in episode i .

TABLE I: Results show the increase of SIN (in %) under different paradigms of interaction on R2RIE-CE benchmark, with localization threshold $\tau_l = 1$, weighting factor $\lambda = 0.01$ from step $p = 4$ onwards. The primary metric **SIN** is highlighted. Under the “No Interaction” column, we report the SR, SPL metrics of the BEVBert policy[13], also showing the Success Rate Upper Bound ($\overline{\text{SR}}$). For *I2EDL*, we set detection threshold $\tau_d = 0.6$. Error type based on R2RIE-CE Val Unseen Dataset.

Error type	No interaction			Random Interaction				Always Ask				I2EDL			
	SR \uparrow	SPL \uparrow	$\overline{\text{SR}}$ \uparrow	SIN \uparrow	MIN \downarrow	SR \uparrow	SPL \uparrow	SIN \uparrow	MIN \downarrow	SR \uparrow	SPL \uparrow	SIN \uparrow	MIN \downarrow	SR \uparrow	SPL \uparrow
Direction	53.4	43.5	58.5	52.9	1.82	53.6	43.6	52.8	3.64	54.3	44.0	53.2	0.50	53.4	43.5
Room*	58.1	48.6	60.4	57.1	1.81	57.9	48.4	56.8	3.62	58.4	48.9	58.1	0.79	58.4	48.8
Object*	56.1	46.1	58.7	55.8	1.75	56.6	46.4	55.2	3.53	56.7	46.7	56.1	0.70	56.3	46.3
Room&Object *	57.3	46.9	61.1	56.9	1.86	57.4	47.3	57.3	3.75	58.4	47.8	58.3	1.15	58.5	47.7
All *	52.4	42.6	61.9	53.3	1.97	53.8	43.5	53.2	3.95	54.1	43.8	53.8	1.37	54.0	43.4
Avg.	55.4	45.5	60.1	55.2	1.84	55.9	45.8	55.0	3.70	56.4	46.2	55.9	0.90	56.1	46.0

Baseline. As *IVLN-CE* is a novel task, there exists no baselines. We thus compare our method *I2EDL* with a “Random Interaction” and an “Always Ask” baseline. Specifically, the “Random Interaction” for every episode i and for every step, randomly predicts if instruction Υ_i contains errors. If the detection is positive (*i.e.*, the instruction contains an error), we then randomly predict a token ℓ_i^j , where $j = \text{rand}(0, \text{len}(\Upsilon_i) - 1)$ and *rand* returns a random number between the arguments. The “Always Ask” baseline prompts an interaction at every step from step $p = 4$ onwards. At each interaction, it randomly predicts the erroneous token in the same way as the “Random Interaction” baseline.

Success Rate and interaction paradigm. In Tab. I, under the “No interaction” column, we report the SR and SPL metrics of the current state-of-the-art method [13] for *VLN-CE*, operating under the different benchmarks of R2RIE-CE [7], thus establishing the lower bound of performances of the agent without interaction. In the *IVLN-CE* scenario, agents should minimize the number of interactions with the human while maximizing their effectiveness. Thus, in Tab. I we report under the “Random Interaction” column the *Success weighted by Interaction Number* (SIN), *Mean Interaction Number* (MIN), SR and SPL scores. We can see that the “Random Interaction”, without being able to distinguish correct instruction from instruction with errors, performs an average of ~ 2 interactions per episode, thus annoying users with unnecessary requests. This behaviour is also reflected in the SIN metric, which penalizes unnecessary interactions. Such behaviour is even more evident with the “Always Ask” baseline. By constantly interacting with the user, the SR is high, however, achieved by an annoyingly large number of user-agent interactions. Finally, in the last columns of Tab. I we present our method *I2EDL*, checking instruction errors from step $p = 4$ onwards, with a detection threshold $\tau_d = 0.6$ and a localization threshold $\tau_l = 1$ (*i.e.*, the predicted token position should differ at most for 1 token). First of all, we note that *I2EDL* has a much higher SIN than the “Random Interaction”, meaning that our method is able to detect instruction errors and localize them more precisely, thus maximizing the effectiveness of the interactions. This is also reflected under the SR column, in which, apart from the *Direction* error benchmark, *I2EDL* has an equal or better SR performance, while halving the average MIN (**0.90** vs 1.84) and scoring consistently lower in terms of SIN. Compared to

the “Always Ask” baseline, *I2EDL* has a higher SIN (**55.9** vs 55.0) while having an extremely low MIN score of **0.90** vs 3.70. Notably, as also reported by [7], *I2EDL* has the lowest results on the *Direction* error benchmark, indicating the challenging of R2RIE-CE.

What is the SR upper bound on R2RIE-CE? In this experiment, we want to establish the Success Rate upper bound ($\overline{\text{SR}}$) that agents can reach in R2RIE-CE, *i.e.*, simulating a perfect agent that is not affected by instruction perturbations. Indeed, ideal agents, as humans, are capable of identifying instruction errors, reasoning and automatically recovering from them. To do this, for each perturbed episode i in \mathcal{E}_p , we substitute the associated perturbed instruction with the correct one, thus giving the correct instruction from the beginning of the episode. Note that we do not change instructions for correct episode i in \mathcal{E}_c . We report the results in Tab. I, under the $\overline{\text{SR}}$ metric. As we can see, the biggest increment is in the *All* benchmark (52.4 vs 61.9), in which the wrong episode $i \in \mathcal{E}_p$ contains three errors. Notably, the second biggest increment is in the *Direction* benchmark, which brings the SR from 53.4 to 58.5. Overall, on average, we have an 8.48% increment from 55.4 to 60.1.

Can agents recover from instruction errors? Here, we want to show the capability of the agent to recover from different instruction errors at different steps. Thus, we want to simulate the following interaction from the human: “Sorry, the instruction I gave you before is wrong. This is the correct one: <instruction>” where <instruction> is the correct instruction. Specifically, for every perturbed episode i in \mathcal{E}_p , we let the agent navigate with the perturbed instruction for t -steps before providing the correct instruction at step $t + 1$. As done before, it is important to note that instructions for correct episodes i in \mathcal{E}_c remain unchanged. We show the results in Fig. 3, where we plot the SR at different t -steps for each benchmark present in R2RIE-CE. As we can see, the *Direction* and *All* benchmarks exhibit the most significant decrease in SR. Note that, *Direction* and *All* have up to one and three errors per instruction, respectively. Particularly for the *Direction* and *All* benchmarks, early error detection is crucial since these errors have a strong impact on the SR.

How do SIN and SR evolve over steps? In this experiment, we want to analyze how SR and SIN evolve across steps using different interaction paradigms. In Fig. 2, we thus show the SR values (solid lines) and SIN values (dashed lines) for

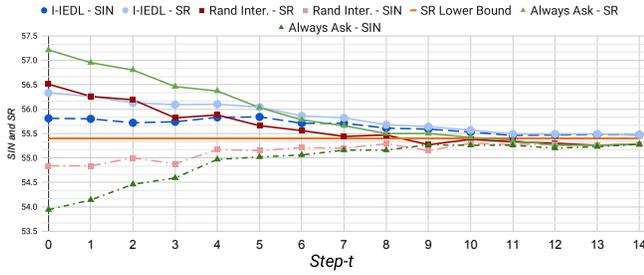


Fig. 2: SR and SIN plotted at different step- t for localization threshold $\tau_l = 1$. Specifically, dashed lines indicate value for the SIN metric, while solid lines for SR. The “Always Ask” baseline always interacts with the user from step- t onwards.

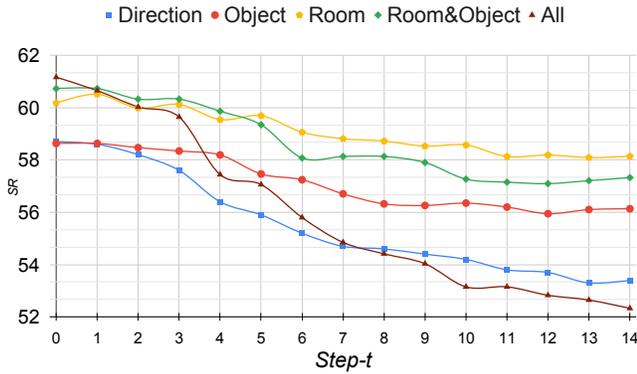


Fig. 3: Success Rate upper bound (\overline{SR}) at different step- t .

different interaction paradigms called for different values step- t . Specifically, (i) “Random interaction” and (ii) “Always Ask”: as described in Sec. IV; and (iii) I2EDL: as described in Sec. IV, in which we set detection threshold $\tau_d = 0.6$. For all the interaction paradigms, $\tau_l = 1$. As we can see from Fig 2, the policy “Always Ask” has the highest SR, since it continuously ask the user. On the other hand, this is very inconvenient, since humans do not want to be disturbed constantly. This behaviour is indeed reflected by our SIN metric, in which “Always Ask” has the lowest SIN values. This behaviour is also reflected in the “Random interaction”. Finally, we can see that our I2EDL has the best compromise between SR and the number of interactions, behaviour that is correctly modelled by the SIN metric.

Conclusions. We presented a novel task, IVLN-CE, which enables interaction between an embodied agent and the human user to correct instruction errors while navigating to a goal described by textual navigation instructions. We proposed an effective baseline, I2EDL, to perform error detection and localization in an online fashion. Compared to baselines, we showed that our proposed I2EDL is generally more effective in improving navigation performance when erroneous instructions are given, while lowering the interaction load. Future works will investigate multi-modal interaction via images and text, with a thorough study on the user’s cognitive load.

REFERENCES

- [1] F. Taioli, F. Cunico, F. Girella, R. Bologna, A. Farinelli, and M. Cristani, “Language-Enhanced RNR-Map: Querying Renderable Neural Radiance Field Maps with Natural Language,” in *ICCVW*, Oct. 2023. [Online]. Available: <http://dx.doi.org/10.1109/iccvw60793.2023.00504>
- [2] P. Anderson, Q. Wu, D. Teney, J. Bruce, M. Johnson, N. Sunderhauf, I. Reid, S. Gould, and A. van den Hengel, “Vision-and-Language Navigation: Interpreting Visually-Grounded Navigation Instructions in Real Environments,” in *CVPR*, Jun. 2018. [Online]. Available: <http://dx.doi.org/10.1109/cvpr.2018.00387>
- [3] J. Krantz, E. Wijmans, A. Majumdar, D. Batra, and S. Lee, *Beyond the Nav-Graph: Vision-and-Language Navigation in Continuous Environments*. ECCV, 2020, p. 104–120. [Online]. Available: http://dx.doi.org/10.1007/978-3-030-58604-1_7
- [4] R. Lloyd, “Cognitive maps: Encoding and decoding information,” *AAG*, vol. 79, no. 1, 1989.
- [5] L. S. Liben, “Conceptual issues in the development of spatial cognition,” in *Spatial cognition*. Psychology Press, 2022, pp. 167–194.
- [6] H. Couclelis, “Verbal directions for way-finding: Space, cognition, and language,” in *The construction of cognitive maps*. Springer, 1996.
- [7] F. Taioli, S. Rosa, A. Castellini, L. Natale, A. D. Bue, A. Farinelli, M. Cristani, and Y. Wang, “Mind the Error! Detection and Localization of Instruction Errors in Vision-and-Language Navigation,” 2024. [Online]. Available: <https://arxiv.org/abs/2403.10700>
- [8] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niebner, M. Savva, S. Song, A. Zeng, and Y. Zhang, “Matterport3D: Learning from RGB-D Data in Indoor Environments,” in *3DV*, oct 2017, pp. 667–676. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/3DV.2017.00081>
- [9] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik, D. Parikh, and D. Batra, “Habitat: A Platform for Embodied AI Research,” in *ICCV*, Oct. 2019. [Online]. Available: <http://dx.doi.org/10.1109/iccv.2019.00943>
- [10] S. Chen, P.-L. Gudur, C. Schmid, and I. Laptev, “History Aware Multimodal Transformer for Vision-and-Language Navigation,” in *NeurIPS*, vol. 34, 2021, pp. 5834–5847. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2021/file/2e5c2cb8d13e8fba78d95211440ba326-Paper.pdf
- [11] Y. Hong, Z. Wang, Q. Wu, and S. Gould, “Bridging the Gap Between Learning in Discrete and Continuous Environments for Vision-and-Language Navigation,” in *CVPR*, Jun. 2022. [Online]. Available: <http://dx.doi.org/10.1109/cvpr52688.2022.01500>
- [12] D. An, H. Wang, W. Wang, Z. Wang, Y. Huang, K. He, and L. Wang, “ETPNV: Evolving Topological Planning for Vision-Language Navigation in Continuous Environments,” *arXiv arXiv:2304.03047*, 2023. [Online]. Available: <https://doi.org/10.48550/arXiv.2304.03047>
- [13] D. An, Y. Qi, Y. Li, Y. Huang, L. Wang, T. Tan, and J. Shao, “BEVBert: Multimodal Map Pre-training for Language-guided Navigation,” *ICCV*, 2023. [Online]. Available: <https://doi.org/10.48550/arXiv.2212.04385>
- [14] W. Zhu, Y. Qi, P. Narayana, K. Sone, S. Basu, X. Wang, Q. Wu, M. Eckstein, and W. Y. Wang, “Diagnosing Vision-and-Language Navigation: What Really Matters,” in *NAACL*, Seattle, United States, Jul. 2022, pp. 5981–5993. [Online]. Available: <https://aclanthology.org/2022.naacl-main.438>
- [15] M. Hahn, A. Raj, and J. M. Rehg, “Which way is ‘right’?: Uncovering limitations of Vision-and-Language Navigation models,” in *AAMAS*, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:253381693>
- [16] Y. Zhang, H. Tan, and M. Bansal, “Diagnosing the Environment Bias in Vision-and-Language Navigation,” in *IJCAI-PRICAI*, Jul. 2020. [Online]. Available: <http://dx.doi.org/10.24963/ijcai.2020/124>
- [17] J. Thomason, M. Murray, M. Cakmak, and L. Zettlemoyer, “Vision-and-Dialog Navigation,” in *CoRL*, ser. Proceedings of Machine Learning Research, vol. 100. PMLR, 30 Oct–01 Nov 2020, pp. 394–406. [Online]. Available: <https://proceedings.mlr.press/v100/thomason20a.html>
- [18] H. Roman Roman, Y. Bisk, J. Thomason, A. Celikyilmaz, and J. Gao, “RMM: A Recursive Mental Model for Dialogue Navigation,” in *EMNLP*. Association for Computational Linguistics, 2020. [Online]. Available: <http://dx.doi.org/10.18653/v1/2020.findings-emnlp.157>
- [19] S. Banerjee, J. Thomason, and J. Corso, “The RobotSlang Benchmark: Dialog-guided Robot Localization and Navigation,” in *CoRL*, ser. Proceedings of Machine Learning Research, J. Kober, F. Ramos, and C. Tomlin, Eds., vol. 155. PMLR, 16–18 Nov 2021, pp. 1384–1393. [Online]. Available: <https://proceedings.mlr.press/v155/banerjee21a.html>

- [20] A. Shrivastava, K. Gopalakrishnan, Y. Liu, R. Piramuthu, G. Tur, D. Parikh, and D. Hakkani-Tur, "VISITRON: Visual Semantics-Aligned Interactively Trained Object-Navigator," 2022. [Online]. Available: <http://dx.doi.org/10.18653/v1/2022.findings-acl.157>
- [21] Y. Zhu, Y. Weng, F. Zhu, X. Liang, Q. Ye, Y. Lu, and J. Jiao, "Self-Motivated Communication Agent for Real-World Vision-Dialog Navigation," in *ICCV*, Oct. 2021. [Online]. Available: <http://dx.doi.org/10.1109/iccv48922.2021.00162>
- [22] T.-C. Chi, M. Shen, M. Eric, S. Kim, and D. Hakkani-tur, "Just Ask: An Interactive Learning Framework for Vision and Language Navigation," vol. 34, no. 03, p. 2459–2466, Apr. 2020. [Online]. Available: <http://dx.doi.org/10.1609/aaai.v34i03.5627>
- [23] K. Nguyen, D. Dey, and B. Brockett, Chrnd Dolan, "Vision-Based Navigation With Language-Based Assistance via Imitation Learning With Indirect Intervention," in *CVPR*, Jun. 2019. [Online]. Available: <http://dx.doi.org/10.1109/cvpr.2019.01281>
- [24] S. Paul, A. K. Roy-Chowdhury, and A. Cherian, "AVLEN: Audio-Visual-Language Embodied Navigation in 3D Environments," 2022. [Online]. Available: <https://doi.org/10.48550/arXiv.2210.07940>
- [25] K. Nguyen and H. Daumé III, "Help, Anna! Visual Navigation with Natural Multimodal Assistance via Retrospective Curiosity-Encouraging Imitation Learning," in *EMNLP-IJCNLP*, 2019. [Online]. Available: <http://dx.doi.org/10.18653/v1/d19-1063>
- [26] P. Anderson, A. Chang, D. S. Chaplot, A. Dosovitskiy, S. Gupta, V. Koltun, J. Kosecka, J. Malik, R. Mottaghi, M. Savva *et al.*, "On evaluation of embodied navigation agents," *arXiv preprint arXiv:1807.06757*, 2018. [Online]. Available: <https://doi.org/10.48550/arXiv.1807.06757>