

Less is More: Generating Grounded Navigation Instructions from Landmarks

Su Wang Ceslee Montgomery Jordi Orbay Vighnesh Birodkar Aleksandra Faust
 Izzeddin Gur Natasha Jaques Austin Waters Jason Baldridge Peter Anderson
 Google Research


Abstract

We study the automatic generation of navigation instructions from 360° images captured on indoor routes. Existing generators suffer from poor visual grounding, causing them to rely on language priors and hallucinate objects. Our MARKY-MT5 system addresses this by focusing on visual landmarks; it comprises a first stage landmark detector and a second stage generator—a multimodal, multilingual, multitask encoder-decoder. To train it, we bootstrap grounded landmark annotations on top of the Room-across-Room (RxR) dataset. Using text parsers, weak supervision from RxR’s pose traces, and a multilingual image-text encoder trained on 1.8b images, we identify 971k English, Hindi and Telugu landmark descriptions and ground them to specific regions in panoramas. On Room-to-Room, human wayfinders obtain success rates (SR) of 71% following MARKY-MT5’s instructions, just shy of their 75% SR following human instructions—and well above SRs with other generators. Evaluations on RxR’s longer, diverse paths obtain 61-64% SRs on three languages. Generating such high-quality navigation instructions in novel environments is a step towards conversational navigation tools and could facilitate large-scale training of instruction-following agents.

1. Introduction

Wayfinding—navigating to a destination—is an everyday task. We study the automatic generation of navigation instructions that effectively guide *people*. Template-based language generators that use cardinal directions and street names are commonly used in outdoor mapping applications, and some more flexible generation approaches rely on databases containing information about maps, roads and landmarks [17, 50, 51]. In contrast, instructions for *indoor* wayfinding require egocentric movement guidance and reference to the visual environment (e.g. notable objects).

Systems for generating indoor wayfinding instructions assume access to pre-existing floorplans and landmark databases [42], but recent work attempts to generate novel



Generated Instruction:

You are facing towards the commode. Turn right and exit the washroom. Turn right and walk straight till you reach the white cabinet in the front. There is an arch in the front. Enter inside the arch. Turn right and walk towards the sofa. Turn left and walk straight till you reach the arch in the front. There is a round table with four chairs towards your left side. You have reached your point.

Figure 1. We generate grounded navigation instructions from a sequence of 360° images captured along a route in a previously unseen building. Our two-stage approach first detects landmarks and then generates instructions conditioned on these landmarks.

instructions directly from visual inputs [22, 39, 58]. Progress toward this goal will enable navigation aids that are conversational rather than map-based—and it could provide a virtually unlimited supply of high-quality synthetic navigation instructions for training instruction-following robots. Describing navigation paths is also a key capability for human-robot communication, equipping robots to answer questions such as *where did you go?* or *where should I meet you?*.

We seek to generate accurate and fluent navigation instructions—in multiple languages—directly from visual representations and actions taken to traverse a path. Previous work assumed that the input to the instruction generator is a sequence of 360° panoramic (henceforth, *pano*) images captured at intervals on a path, typically training on instructions from Room-to-Room (R2R) [5] using Matterport3D environments [9]. These models’ instructions have proven

valuable as additional training data for vision-and-language navigation (VLN) agents [22]. However, *people* struggle to follow them [64]: human wayfinding success rates on R2R are 36% for Speaker-Follower [22] and 42% for EnvDrop [58] in unseen environments. The generated text is stylistically correct, but frequently references non-existent objects and confuses spatial terms such as left and right.

A challenge for visually-oriented instruction generators is dealing with irrelevant visual inputs. In many other image-to-text generation tasks (e.g., image captioning), much of the visual information in the input is reflected in the output text. This is not the case when generating navigation instructions. Human annotators look at less than 30% of the environment [36], and the instructions reference only a fraction of the objects that they look at. This makes learning a precise mapping between visual inputs and text outputs much harder. Perversely, access to more information can degrade performance [15], as models happily learn spurious correlations that cause hallucinations during inference.

To solve this, we exploit the spatiotemporal grounding in the Room-across-Room (RxR) dataset [36]. Instead of *writing* instructions, RxR annotators *spoke* while traversing paths. Every RxR instruction thus comes with *pose traces* that align the words spoken (and later transcribed) with what annotators were looking at. We use these pose traces and instructions to derive a new *silver* annotated dataset¹ that contains bounding boxes over visual landmarks combined with their multilingual descriptions (English, Hindi and Telugu). Specifically, we bootstrap landmark annotations using text parsers to identify landmark phrases in instructions. We then use powerful image-text co-embedding models [32] combined with weak supervision from pose traces to ground those landmarks in the environment.

Our two-stage MARKY-MT5 (**landmark** and **multilingual T5** [62]) system enhances instruction generation by improving how visual landmarks are selected and mentioned. Given a path-connected sequence of panoramic views, the first stage **landmark detector** infers a sequence of landmarks that a person might select for describing the path. E.g., in Fig. 1 eight landmarks are selected, each represented by an image. This sequence, plus interleaved descriptions of navigation actions, is passed to the second stage **instruction generator** – a multimodal extension of the multilingual T5 (mT5) model [62] similar to VL-T5 [13] – to produce the instruction in Fig. 1.

In human wayfinding experiments on R2R paths, MARKY-MT5 trained with silver landmarks (a subset of the visual inputs from the full environment) almost eliminates the gap between model-generated and human-written instructions – achieving a 71% success rate (SR) vs. 75% for human instructions, 42% for previous models, and 58%

for our model trained on full 360° panos. When it comes to selecting visual inputs for the generator, *less is more*. On the more challenging RxR paths, human wayfinders obtain a 62% SR using MARKY-MT5 vs. 78% for human instructions. We release our silver landmark data and over one million navigation instructions generated by MARKY-MT5, as data augmentation for training VLN agents.²

2. Related Work

Wayfinding with landmarks. We wish to produce instructions that *people* can follow, and are inspired by research on the importance of landmarks for human navigation [8, 18, 21, 63]. Landmarks are not just spatial features – they encode a relation between the feature (e.g. object) itself, its nearby environment, and wayfinders’ point of view [7]. Our landmark detector is trained on data bootstrapped from RxR’s human-referenced landmarks. This allows our approach to exploit such features of landmark saliency, without designing for them explicitly (as in [19, 27]).

Hong *et al.* [29] show that modeling relationships among the scene, its objects, and directional clues is effective for improving VLN wayfinding performance. This points to a potentially virtuous cycle between agents for wayfinding and those for guiding—or, better, using such landmark understanding for both capabilities within individual agents.

Navigation instruction generation. Previous work on generating synthetic VLN instructions employs the Speaker-Follower framework [23, 58]: A Speaker model learns from R2R annotations (English only) to produce instructions conditioned on sequences of path panoramas, while a Follower model learns to wayfind (i.e. construct paths) conditioned on human instructions and the same visual inputs. The Speaker’s output can be used as augmented data for training the Follower, and the Speaker is used to rerank paths generated by the Follower during inference (a form of pragmatic reasoning). These models indiscriminately use entire panos as visual context, whereas we select key visual landmarks from each pano for the generator to talk about. We build on the multitask, multilingual T5 model architecture [48, 62], a unified text-to-text framework that enables transfer learning by mixing many NLP tasks simultaneously. This additionally allows us to explore pretraining tasks, including multimodal tasks like image captioning, to improve generalization on unseen environments.

Agarwal *et al.* [1] previously proposed a landmark-based generator, but relied on RL training rather than silver data to induce landmark groundings. Pashevich *et al.* [45] use synthetic instructions like *goto bed pickup cellphone* as an additional source for training VLN agents for the ALFRED benchmark [54]. These are akin to the simple directional expressions we use in our multimodal encoder, but they are

¹The term *silver data* refers to high-quality annotations—not created by people—that are derived by combining models and constraints [26, 49, 61].

²github.com/google-research-datasets/RxR/tree/main/marky-mT5

used both for data augmentation and as an additional decoding task. They seek to optimize VLN agent performance, whereas we seek to produce instructions that can be followed by people.

Kojima *et al.* [35] explore collaborative, human-agent instruction generation in the CEREALBAR game [57]. They define a human-in-the-loop instruction generation framework, where the generator is iteratively improved with signals collected while *interacting* with people. The instructions cover both navigation and game tactics.

Their multimodal→text generator is a blackbox, whereas our approach includes an interpretable intermediate representation based on selected visual landmarks (compared to a formal abstraction, e.g. [14]). This work complements ours—which only uses static human annotations—and suggests future interactive settings that could allow our instruction generators to adapt to human wayfinders.

Multimodal generation. Our use of landmarks and fine-grained connections between language and visual elements has parallels in image captioning. In particular, Pont-Tuset *et al.* [46] show that controlled image captioning using mouse traces produces better image descriptions. Zhou *et al.* [66] use a pre-trained text-image matcher to learn keyword-bbox alignments, which are used to condition their caption generator (similar to our detection and use of visual landmarks). Huber *et al.* [30] provide a two-stage image-grounded dialogue model that extracts sentiment information from images to imbue emotion into generated texts. Specialized datasets, akin to our silver landmark data, are used to train an image feature extractor to improve scene understanding and handling of sentiment and facial features.

Our generator itself falls within the growing collection of approaches that encode multimodal inputs and decode to text, including VL-T5 [13], MAnTiS [55] and SimVLM [60], often in a multitask setting. Our inputs are different in that we encode multiple, connected images that are interleaved with action descriptions. Our model outputs are evaluated with respect to downstream task performance using human evaluators, rather than automatic metrics and downstream tasks based on the learned representations.

3. Bootstrapping a Landmark Dataset

MARKY-MT5’s first stage identifies visual landmarks as input to a second stage instruction generator. This requires training data of navigation instructions with landmark annotations. Here, a *landmark annotation* is an instruction text span that mentions an object (e.g., *white couch*), and the aligned perspective image containing that object (see Fig. 2). Unfortunately, navigation instruction datasets with annotated landmarks do not exist.³ Given how costly man-

³While Matterport3D contains 3D bounding box annotations for many objects (potential landmarks), these are not aligned to navigation instructions (which typically refer to a much broader set of landmarks than these).




Figure 2. Bootstrapped landmarks illustrated as image-text pairs.

ual annotation would be at scale, we instead opt to bootstrap *silver data*, a term used in natural language processing when high-quality annotations – not created by people – are derived by combining models and constraints [26, 49, 61].

RxR. To automatically identify and annotate landmarks, we use *Room-across-Room* (RxR) [36]. RxR builds on the core task formulation defined by *Room-to-Room* (R2R) [5]. Both are set in Matterport3D environments [9], but RxR is multilingual (English, Hindi, Telugu), its paths are longer and more varied, and it has roughly six times more instructions. Importantly, while both R2R and RxR include navigation paths, RxR includes two key additional annotations:

- *Pose traces.* A pose trace records an annotator’s virtual pose as they move along a navigation path while recording spoken navigation instructions for others to follow. From a pose trace we can render a timestamped video of everything the annotator looked at.
- *Text timestamps.* Every word in RxR’s text instructions is timestamped, and thus aligned to the pose trace video. This provides a noisy, approximate grounding of words in the navigation instructions to video frames.

Using RxR, we construct landmark annotations in three steps: (1) *extracting landmark phrases*; (2) *grounding landmark phrases to pose trace frames*; and (3) *refining selected frames to better align with the landmark object*. Text timestamps provide a key source of weak supervision for step 2, since an annotator would generally utter “*white couch*” soon after seeing the couch. However, text timestamps alone are insufficient for grounding to pose trace frames as annotators may talk about a landmark without looking at it.

Extracting landmark phrases. The first step to create silver landmark data is extracting temporally-ordered lists of landmark phrases – noun phrases describing objects in the environment – from RxR’s human instructions. For this,

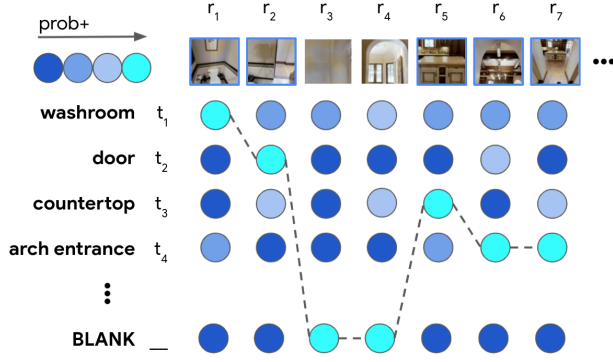


Figure 3. Example CTC alignment between pose trace frames (top) and landmark phrases (left). Every landmark is aligned to a contiguous subsequence of frames. The BLANK label is used for frames that do not contain a landmark.

we use an mBERT-based [16] dependency parser trained on multilingual Wikipedia data [25] to identify all entity mentions, including their parts-of-speech, which we consolidate with their non-clausal dependents into a single text span (see supplementary for further details). To improve results, in each language we manually curate a stop list of common landmark phrases that are discarded as too difficult to accurately localize (e.g., ‘the room’, ‘your destination’). The output is a text span sequence $\mathbf{t} = [t_1, t_2, \dots, t_m]$, e.g.

*You are standing in front of a **brown chair**. Take a left to enter the **bathroom**, you will see a **sink** in front of you. Now take a step to the right at stop at the **foot mat**, you will have reached your destination.*
 \rightarrow [**brown chair**, **bathroom**, **sink**, **foot mat**].

In total, we extract 971k landmark phrases from the 102k instructions in the RxR Train, Val-Seen and Val-Unseen splits (9.55 landmarks per instruction on average). There are 84,237 unique phrases; the percentage of phrases that appear more than 3/5/10 times is 20.8%/13.1%/7.4%.

Landmark grounding matrix. For an instruction, to ground landmark phrases $\mathbf{t}=[t_1, t_2, \dots, t_m]$ to pose trace video frames $\mathbf{r}=[r_1, r_2, \dots, r_n]$, we construct an $m \times n$ logit matrix A , where $A_{i,j}$ represents the compatibility of landmark phrase t_i and frame r_j . Logits are computed by combining signals from MURAL-large [32] – a high-performing multilingual, multimodal dual encoder trained on a mixture of 1.8b noisy image-text pairs and 6b translation pairs – and the RxR text timestamps, i.e.:

$$A_{i,j} = X(t_i) \cdot Y(r_j) - \lambda(T(t_i) - T(r_j))^2 \quad (1)$$

where X computes a MURAL text embedding, Y computes a MURAL image embedding, T returns a timestamp in seconds, and λ is a weighting factor set to 1.

To finetune MURAL on RxR, we add an additional trainable layer to each pretrained and frozen encoder and

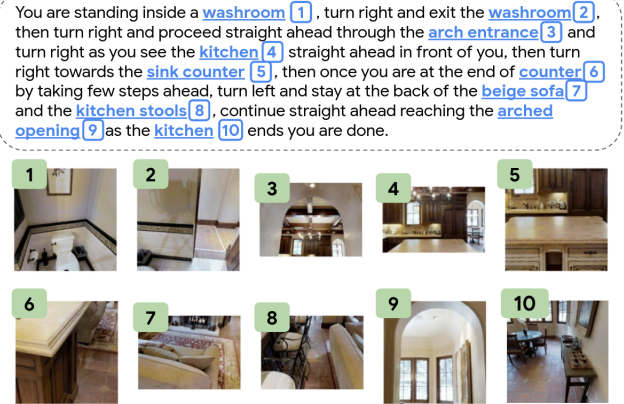


Figure 4. A full bootstrapped landmark annotation for a navigation instruction from RxR Val-Unseen. Nine of ten landmarks (1-9) are strongly aligned. Landmark 10 *kitchen* is misaligned (despite the visual resemblance), perhaps due to a violation of our assumption that annotators mention landmarks in their observation order – here, the mention of *kitchen* harks back to an earlier observation.

use *Connectionist Temporal Classification* (CTC) loss [24]. Our motivation is similar to uses of CTC in speech recognition and optical character recognition [20], where the sequence of output labels is known, but the alignment to the inputs is not. We treat the frame sequence $\mathbf{r}=[r_1, r_2, \dots, r_n]$ as the input. The set of output labels $\{t_1, t_2, \dots, t_m, \text{BLANK}\}$ includes the landmark phrases in the instruction plus a BLANK label (to accommodate frames without a landmark). The ordered landmark phrases $\mathbf{t}=[t_1, t_2, \dots, t_m]$ provide the target output sequence, i.e., we assume that the order in which annotators mention landmarks is the order in which they are observed. CTC’s loss maximizes the probability that predicting a landmark phrase for each frame (or BLANK) produces a valid output sequence:

$$p(\mathbf{t} | \mathbf{r}) = \sum_{A^{(\mathbf{r}, \mathbf{t})}} \prod_{j=1}^n p(A_{i,j} | r_j) \quad (2)$$

where $A^{(\mathbf{r}, \mathbf{t})}$ marginalizes over all valid alignments, and the right-hand side item computes the probability of single alignment. Intuitively, CTC performs training by encouraging each frame to be assigned to a landmark phrase such that the order of landmarks in the instruction is preserved. A valid CTC alignment is demonstrated in Figure 3.

Decoding landmark alignments. After finetuning, we perform greedy CTC decoding on the logit matrix A which aligns each landmark phrase with a sequence of one or more frames. For each landmark phrase t_i we then find the best single alignment by selecting the frame in that sequence that assigns t_i the highest probability. This generates image-text pairs where each text span is a landmark phrase from an instruction, and each image is a perspective projection from a 360° pano with known heading, pitch, and horizontal and vertical field of view (derived from the original pose trace).

Evaluation. As the text-image alignment task above is weakly-supervised (i.e., we know a landmark phrase sequence and a pose trace video are paired, but we do not know how landmark phrases are aligned with video frames), we also perform a small-scale human evaluation of alignment quality. We manually annotated 100 instructions from RxR Val-Unseen, evaluating our approach in various conditions using alignment precision. We observe: 1) gains with adding timestamp-based bias (i.e. the squared term in Eq. (2)): 49.6 w/o bias vs. 56.5 w/ bias; 2) The bias alone (i.e. no input from MURAL) does not work: 23.4 precision.

Landmark refinement. We observe that the pose trace frames selected in the previous steps typically contain the landmark phrase, but it is not always centered. This is because annotators do not necessarily need to look directly at a landmark to notice it. To further refine the landmark grounding, we generate a large number of alternative groundings (defined by heading, pitch and field of view) that overlap with the selected frame. The final silver landmark image for each landmark phrase is the grounding among these candidates with the highest MURAL alignment score with the landmark phrase. This results in very high quality landmark annotations, as shown in Fig. 4 for a complete set of landmark annotations for one instruction.

4. Landmark Detection

Using navigation paths from RxR and silver landmark annotations (Sec. 3), we train a *landmark detector* that selects visual landmarks to be mentioned. The input to the model is a sequence of 360° panos $I_{1:T}$ captured at regular intervals along a route, along with the inbound and outbound direction. The model detects landmarks in 360° images as an object detection problem, using a CenterNet [65] object detector. The detector is intended for identifying landmarks in unseen environments. Although not all landmarks are objects in the traditional sense (e.g. *open area*, *laundry room*), previous work demonstrates that object detectors can learn to successfully detect and featurize amorphous image regions such as *sky* and *water* [4].

Input format. Given a sequence of panoramic images $I_{1:T}$, each pano I_t is input to CenterNet in equirectangular format. Though equirectangular images suffer from distortion, most landmarks are located close to the horizon, where equirectangular distortion is minimal. Previous work on landmark salience for human navigation suggests that landmarks located closer to the next route segment are more important than landmarks further away (e.g. behind the observer) [7, 34]. We therefore provide route context to the model, by rotating each input pano so that the direction to the next pano (the *outbound direction*) is centered in the image. Since the direction faced on arrival from the previous pano (the *inbound direction*) is also important, we draw colored pixel blocks on the input image indicating both to the

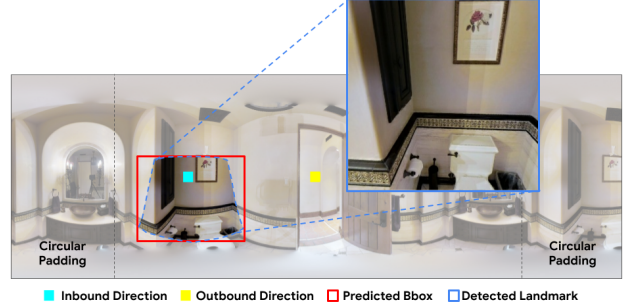


Figure 5. Pano input format for the landmark detector. Both the inbound and outbound directions are drawn on the image to provide route context. The detector outputs bounding boxes in the equirectangular image, which are converted to perspective projections of the detected landmark(s).

model (see Fig. 5). To avoid edge discontinuities, we also add 90° of circular padding to each side of the pano. Images are prepared consistently for both training and inference.

Training. As outlined in Sec. 3, each landmark image is a perspective projection from a 360° pano. To obtain training targets we map these projections to their minimum enclosing bounding box in the equirectangular pano image. We initialize the model with an Hourglass-104 backbone [40, 43] pretrained on the COCO 2017 dataset [41], and train for 20k iterations on silver landmarks from RxR training set using a batch size of 128, an equirectangular image size of 512×256 , and a single landmark object class.

Inference. During inference, we pool together the top 3 detections from each pano in the path, and return the top T detections from this pool as the detected landmarks (where T is the path length).⁴ These values (at most 3 landmarks from a single pano and one landmark per pano on average) closely correspond to averages in the training set.

5. Instruction Generation

The instruction generator converts a representation of selected landmarks and the route connecting them into a textual navigation instruction. Our model is based on the mT5 text-to-text encoder-decoder transformer architecture [62], which is itself a multilingual variant of T5 [48] (hence, MARKY-MT5). In all experiments we use the mT5 base size model and the standard mT5 vocabulary which supports up to 101 languages using a SentencePiece [38] model trained on mC4 to encode text as WordPiece [52] tokens.

Input format. We formulate the input to the model by inserting visual representations of selected landmarks into a templatic English text sequence describing the orientation of each landmark and the actions required to traverse the route. Specifically, each visual landmark feature appears in

⁴The 1:1 #landmarks / path length ratio is selected based on BLEU and CIDEr scores: also experimented with 1.2:1, 1.5:1, 1.75:1 and 2:1 ratios.

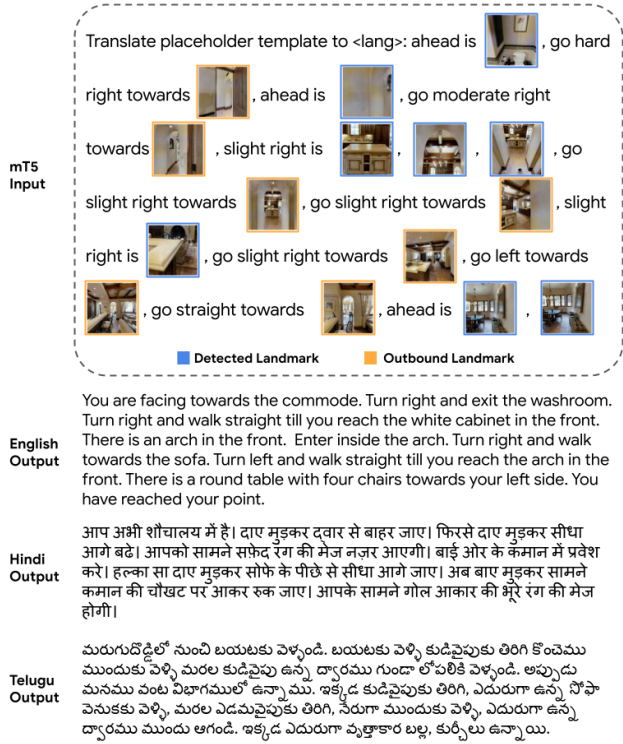


Figure 6. Input template and multilingual outputs from the instruction generator. The input template encapsulates the visual landmarks to be mentioned along the route, along with their relative orientation and the navigation actions connecting them.

order of its appearance on the route, preceded by its relative orientation to an observer, e.g. *ahead is...*, *slight right is...*. In between each landmark, the relative motion required to approach the next landmark is encoded with an action phrase, e.g. *go hard left*, *turn 180* (refer Fig. 6). There are 12 possible relative orientation phrases and 12 possible action phrases, with each one representing an arc of 30° . Using this approach, our instruction generator can be trained or evaluated using any set of landmark images that can be oriented along a 3D route. To identify the desired output language, the input is preceded by the appropriate prompt: ‘Translate placeholder template to English/Hindi/Telugu:’.

Landmarks. To train the model we use the RxR dataset augmented with the silver landmarks described in Sec. 3. Since instructions often mention the same landmark several times (e.g. from Fig. 4 ‘You are standing in a *washroom*, turn right and exit the *washroom...*’), the same landmark image can appear multiple times in succession in the silver data. This leaks information about the surface realization of the instruction which will not be available from our landmark detector. We therefore de-duplicate landmarks by removing any landmark with heading, elevation and field of view within 5° of a previous landmark from the same pano.

In addition to the landmarks provided by silver data or

the landmark detector (*detected landmarks*), at each pano we include an additional *outbound landmark*, defined as the view from the current pano in the direction of the next pano (bordered in orange in Fig. 6). Including outbound landmarks gives a small boost in automatic evaluations. All landmarks are represented using 640-dimension image embeddings from MURAL-large [32].

Rewrite task. To improve the performance of the instruction generator, we propose an auxiliary *rewrite task*. In the rewrite task, we modify the input template during training by replacing the image feature for each silver landmark with the landmark phrase, e.g. ‘commode’. As before, the training target is a full navigation instruction containing all the landmark phrases in context. We hypothesize that this auxiliary task helps train the attention mechanisms in the model to correctly associate input landmark representations with their output realization in the text (noting that this is the only way we make direct use of the landmark phrase annotations from the silver dataset).

Pretraining. We investigate pretraining the instruction generator on an equally weighted mixture of image-caption pairs from the CC3M [53] and CC12M [10] datasets. Since the captions are all in English, we translate them into Hindi and Telugu using an MT service.⁵ We consider it important to introduce both text and visual inputs to our model during pretraining. We therefore formulate the pretraining task as machine translation of image captions using a pivot image [56], e.g. (English + image) \rightarrow Hindi and so on. To encourage the model to make full use of the visual information only the first half of the input caption is provided.

Finetuning and inference. We finetune mT5 on all three languages in the RxR train split using silver landmark inputs for 200k iterations with batch size 128. During inference we use argmax decoding and evaluate using both silver and predicted landmarks from the landmark detector.

6. Experiments

Baselines. We train our multilingual MARKY-mT5 model on RxR and evaluate using paths from both RxR [36] and R2R [6], comparing to the following prior work:

- **SpkFol** and **EnvDrop**: Instruction generators from the Speaker-Follower [22] and Environmental Dropout [58] papers. Both models are LSTM [28] encoder-decoder models that take full panoramic images as input, with each pano represented by 36 image features representing different viewing directions.
- **SpkFol-RxR**: We re-implement SpkFol and optimize hyperparameters for the longer instructions in the RxR dataset, training separate monolingual models for English, Hindi and Telugu using the mT5 vocabulary.

⁵<https://cloud.google.com/translate>

	Model	Landmarks	Training Data	WC	NE ↓	SR ↑	SPL ↑	Quality ↑	Visual Search %		
									Start ↓	Other ↓	Time (s) ↓
R2R (en)	1 SpkFol [22]	Full Panos	R2R	24.6	6.0	42.0	35.8	4.1	39.8	23.6	54.2
	2 EnvDrop [58]	Full Panos	R2R	24.5	6.0	41.7	35.3	4.0	40.7	23.5	54.0
	3 SpkFol-RxR [22]	Full Panos	RxR	61.8	3.9	57.8	48.7	4.2	36.0	23.7	67.5
	4 Marky-mT5	Outbound	RxR	57.5	3.6	64.9	54.1	4.2	36.2	23.8	72.5
	5 Marky-mT5	Predicted	RxR	58.2	2.9	70.8	59.8	4.3	35.5	23.2	70.1
	6 Human	-	-	25.6	2.8	74.9	66.4	4.5	37.8	23.0	52.2

Table 1. R2R Val-Unseen human wayfinding performance ($N = 783$ for each model). Combining the larger RxR dataset with landmark modeling and our bootstrapped landmark dataset, we almost eliminate the gap between model-generated and human-written instructions on paths of R2R-level difficulty – achieving a 70.8% success rate vs. 74.9% for human instructions and 42% for previous models.

	Model	Landmarks	WC	NE ↓	SR ↑	SDTW ↑	NDTW ↑	Quality ↑	Visual Search %		
									Start ↓	Other ↓	Time (s) ↓
RxR (all)	1 Marky-mT5	Outbound	67.3	5.3	53.3	37.4	52.1	4.1	38.9	26.9	146.8
	2 Marky-mT5	Predicted	71.1	4.4	61.5	43.9	56.2	4.2	37.6	26.5	154.4
	3 Marky-mT5	Silver	72.2	4.3	63.3	45.4	57.0	4.3	37.8	26.6	149.8
	4 Human	-	76.2	2.7	78.4	61.3	69.2	4.6	36.5	25.7	147.4

Table 2. RxR Val-Unseen human wayfinding performance ($N = 4,551$ for each model). On longer and more challenging RxR-style paths, human wayfinders obtain a 61.5% success rate using our model vs. 78.4% for human instructions. The higher 63.3% success rate using silver landmarks suggests improved landmark selection is one way to narrow that gap.

To evaluate modeling choices and opportunities, we report results for the following variations of our model:

- **Outbound:** Training and evaluation with outbound landmarks, but no detected landmarks (refer Sec. 5).
- **Predicted:** Our full model, trained on silver landmark data and evaluated using predicted landmarks generated by our landmark detector.
- **Silver:** As above, evaluated with silver landmarks. This allows us to assess a quasi-upperbound on performance on landmarks *derived* from human annotations and does *not* represent deployed system performance.
- **Ablations:** Results without pretraining on CC3M and CC12M, and without the Rewrite auxiliary task.

Human Wayfinding. We consider human evaluations to be essential for this task. Using PanGEA [37], an open-source annotation toolkit for panoramic graph environments, we immerse annotators in a simulated first-person environment backed by the Matterport3D dataset [9] and ask them to follow the provided text navigation instructions. In total, we conduct 20k wayfinding evaluations involving 70 annotators. All annotators are based in India.

Human Wayfinding Metrics. We report the following standard metrics [2, 31] to evaluate the similarity between the route generated by our annotators and the intended route (and hence, the quality of the generated instructions): navigation error (NE ↓) success rate (SR ↑), success weighted by inverse path length (SPL ↑), normalized dynamic time warping (NDTW ↑), and success weighted DTW (SDTW ↑). Arrows indicate improving performance. Following [64], we also report instruction quality assessed by annotators on a 1–5 Likert scale (Quality ↑), and the percentage of the available panoramic visual field that the an-

notator observes at each viewpoint (**Visual Search** ↓), with higher values indicating greater effort spent looking for the correct route or landmarks.⁶ **Time** ↓ represents the average time taken in seconds to complete the task, **WC** is word count and **N** is the number of instructions evaluated.

Automatic Metrics. For model development and ablations, we use automatic evaluation metrics. SPICE [3] is the only metric found to correlate with human wayfinding performance [64]; however, SPICE needs an English dependency parser [33], so we also report BLEU [44] and CIDEr [59]. Further, since SPICE was designed for evaluating individual sentences, not multi-sentence paragraphs, when calculating SPICE we separate both candidate and reference instructions into individual sentences (details in Supplementary). We compute CIDEr and SPICE with official evaluation code from COCO captions [12]. For BLEU, we use the sacre-BLEU implementation [47]. CIDEr and BLEU scores are macro-averaged over languages.

Results. We present human evals on R2R’s Val-Unseen paths in Tab. 1 and RxR’s Val-Unseen paths in Tab. 2; automated evaluations on RxR are in Tab. 3. We align our results discussion to specific questions and themes.

Dataset choice: Training on the larger RxR dataset instead of R2R noticeably improves human wayfinding performance in English, increasing success rate (42.0% → 57.8%) and SPL (35.8% → 48.7%) while lowering navigation error (6.0m → 3.9m) using the same model – see Tab. 1 row 3 vs. row 1. We note that the word count (WC) of models trained on RxR is considerably longer than R2R

⁶**Start** and **Other** refer to first and other viewpoints, reported separately because wayfinders usually look around to orient themselves at the start.

				Val-Seen			Val-Unseen					
Model		Landmark	Rewrite	PT	BLEU	CIDEr	SPICE (en)	BLEU	CIDEr	SPICE (en)	SR \uparrow	NDTW \uparrow
RxR	1 SpkFol-RxR [22]	Full Panos	-	-	5.9	8.9	13.9	5.7	8.4	13.0	29.6	41.6
	2 Marky-mT5	Full Panos	-	-	5.8	8.3	11.8	5.6	9.6	14.0	50.7	60.1
	3 Marky-mT5	Outbound	-	-	6.4	10.7	10.1	6.2	10.0	10.0	53.6	62.9
	4 Marky-mT5	Silver	-	-	11.6	25.2	16.6	11.1	23.3	15.7	55.9	64.1
	5 Marky-mT5	Silver	\checkmark	-	14.0	33.9	17.8	12.9	30.4	16.4	56.3	64.2
	6 Marky-mT5	Silver	\checkmark	\checkmark	14.6	35.0	17.3	13.4	31.7	16.2	56.4	64.2
	7 Marky-mT5	Predicted	\checkmark	\checkmark	6.2	8.4	13.7	5.8	7.5	13.5	55.7	63.3
	8 Human	-	-	-	-	-	-	-	-	-	56.5	62.9

Table 3. Automatic evaluations on RxR. Results are aggregated across English, Hindi and Telugu except for SPICE which is English-only. SR and NDTW score the wayfinding performance of a state-of-the-art VLN agent [11]. PT refers to CC3M/12M PreTraining.

for the same paths; RxR annotations include more detailed descriptions and more state verification [36].

Do landmarks matter? Yes. On R2R, MARKY-MT5 increases success rate (57.8% \rightarrow 70.8%) and SPL (48.7% \rightarrow 59.8%), and lowers navigation error (3.9m \rightarrow 2.9m) compared to prior work without landmarks – represented by SpkFol-RxR trained on the same dataset (Tab. 1, row 5 vs. 3). We attribute our gains primarily to landmark modeling: though MARKY-MT5 uses a transformer architecture and stronger image features than SpkFol-RxR, these changes have little impact without landmarks (Tab. 3, row 2 vs. 1).

How useful are the silver landmarks? On RxR’s challenging paths, training and evaluating with silver landmarks increases success rate (53.3% \rightarrow 63.3%), increases NDTW (52.1% \rightarrow 57.0%), and lowers navigation error (5.3m \rightarrow 4.3m) compared to using only outbound landmarks (Tab. 2, row 3 vs. 1). This confirms our silver landmarks contain useful groundings and using them improves performance.

Can landmark detection be improved? MARKY-MT5 performs better using silver landmarks instead of predicted landmarks, e.g. success rate of 63.3% vs. 61.5% (Tab. 2 row 3 vs. 2), suggesting gains are possible.

Comparison to human instructions: Overall, using a combination of RxR data, silver landmarks and modeling improvements, we almost eliminate the gap between model-generated and human-written instructions on paths of R2R-level difficulty – with a 71% success rate vs. 75% for human instructions and 42% for previous models (Tab. 1). However, on the more challenging RxR-style paths, a gap remains – human wayfinders obtain a 62% success rate using MARKY-MT5 vs. 78% for human instructions (Tab. 2). While humans can still distinguish MARKY-MT5 instructions from human, a state-of-the-art VLN agent [11] cannot – achieving near identical success rates (56.5% vs. 55.7%) and NDTW (62.9% vs. 63.3%) for human and model-generated instructions respectively (Tab. 3 row 8 vs. 7).

Rewrite auxiliary task and CC3M/12M pretraining: Training with the Rewrite task (Tab. 3 row 5 vs. 4) has a large positive effect on Val-Seen and Val-Unseen BLEU (+2.4/+1.8), CIDEr (+8.7/+7.1) and SPICE (+1.2/+0.7). Pretraining (Tab. 3 row 6 vs. 5) improves BLEU (+0.6/+0.5)

and CIDEr (+1.1/+1.3) scores, but not SPICE (-0.5/-0.2).

Trustworthiness of automatic evaluations: We note that replacing silver landmarks with predicted landmarks (Tab. 3 row 7 vs. 6) dramatically reduces BLEU (-8.2/-7.6), CIDEr (-26.6/-24.2) and SPICE (-3.6/-2.7) to the level of previous work. This is completely inconsistent with our human evaluations. We suggest *only using automatic evaluations to compare similar models*, and attribute this reduction in textual similarity to differences in the choice of landmarks.

Diverse generation. An appealing property of our two-stage approach is that diverse instructions can be generated by sampling landmark predictions. Revisiting the Fig. 1 example, sampling the next highest ranked landmark predictions produces the following–quite different–instruction:

”You are standing in a bathroom facing a toilet. You are going to turn to your right and exit the bathroom. You will come into a kitchen area. There will be a marble counter top on your right and to your left a living room. You are going to make a left into the living room. In front of you will be an arched entrance way with a table and chairs. You are to stop right in front that entrance way and you are done.”

7. Conclusion

Supported by a new bootstrapped dataset of 971k grounded landmarks, our MARKY-MT5 model almost eliminates the gap between model and human-written instructions on R2R paths. Generating such high-quality navigation instructions in novel environments is a step towards conversational navigation tools and could enable larger-scale training of instruction-following agents. However, the strength of our approach–focusing on visual landmarks–is also a limitation. MARKY-MT5 is blind to other context when generating, making it susceptible to pragmatic failures, e.g. generating ‘Leave the room’ in a room with multiple exits. Addressing this could lead to further gains.

Acknowledgements. We thank Ming Zhao, Subhashini Venugopalan, and Alex Ku for early discussions; Yinfei Yang, Chao Jia and Aashi Jain for help with MURAL image features; Sebastian Goodman and Beer Changpinyo for assistance with the multi-modal mT5 implementation; and Igor Karpov, Ming Zhao and the Google ML Data Operations team for supporting human evals.

References

- [1] Sanyam Agarwal, Devi Parikh, Dhruv Batra, Peter Anderson, and Stefan Lee. Visual landmark selection for generating grounded and interpretable navigation instructions. In *CVPR workshop on Deep Learning for Semantic Visual Navigation*, 2019. 2
- [2] Peter Anderson, Angel X. Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, and Amir Roshan Zamir. On evaluation of embodied navigation agents. *CoRR*, abs/1807.06757, 2018. 7
- [3] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. SPICE: Semantic Propositional Image Caption Evaluation. In *ECCV*, 2016. 7
- [4] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 2018. 5
- [5] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 3
- [6] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. Vision-and-Language Navigation: Interpreting visually-grounded navigation instructions in real environments. In *CVPR*, 2018. 6
- [7] David Caduff and Sabine Timpf. On the assessment of landmark saliency for human navigation. *Cognitive processing*, 9(4):249–267, 2008. 2, 5
- [8] Edgar Chan, Oliver Baumann, Mark A. Bellgrove, and Jason B. Mattingley. From objects to landmarks: the function of visual location information in spatial navigation. In *Frontiers in Psychology*, 2012. 2
- [9] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3D: Learning from RGB-D data in indoor environments. *3DV*, 2017. 1, 3, 7
- [10] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *CVPR*, 2021. 6
- [11] Shizhe Chen, Pierre-Louis Guhur, Cordelia Schmid, and Ivan Laptev. History aware multimodal transformer for vision-and-language navigation. In *NeurIPS*, 2021. 8
- [12] Xinlei Chen, Tsung-Yi Lin Hao Fang, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollar, and C. Lawrence Zitnick. Microsoft COCO Captions: Data Collection and Evaluation Server. *arXiv preprint arXiv:1504.00325*, 2015. 7
- [13] Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. Unifying vision-and-language tasks via text generation. In *Proceedings of ICML 2021*, 2021. 2, 3
- [14] Andrea F. Daniele, Mohit Bansal, and Matthew R. Walter. Navigational instruction generation as inverse reinforcement learning with neural machine translation. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction, HRI '17*, page 109–118, New York, NY, USA, 2017. Association for Computing Machinery. 3
- [15] Pim de Haan, Dinesh Jayaraman, and Sergey Levine. Causal confusion in imitation learning. In *NeurIPS*, 2019. 2
- [16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, June 2019. 4
- [17] Markus Dräger and Alexander Koller. Generation of landmark-based navigation instructions from open-source data. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 757–766, Avignon, France, Apr. 2012. Association for Computational Linguistics. 1
- [18] Russell A. Epstein and Lindsay K. Vass. Neural systems for landmark-based wayfinding in humans. In *Philos Trans R Soc Lond B Biol Sci*, 2014. 2
- [19] Irene Fellner. Automatic generation of landmark-based indoor routing instructions. 2017. 2
- [20] Xinjie Feng, Hongxun Yao, and Shengping Zhang. Focal ctc loss for chinese optical character recognition on unbalanced datasets. In *Complexity*, 2018. 4
- [21] Patrick Foo, William H. Warren, Andrew Duchon, and Michael J. Tarr. Do humans integrate routes into a cognitive map? map- versus landmark-based navigation of novel shortcuts. In *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 2005. 2
- [22] Daniel Fried, Ronghang Hu, Volkan Cirik, Anna Rohrbach, Jacob Andreas, Louis-Philippe Morency, Taylor Berg-Kirkpatrick, Kate Saenko, Dan Klein, and Trevor Darrell. Speaker-follower models for vision-and-language navigation. In *NeurIPS*, 2018. 1, 2, 6, 7, 8
- [23] Daniel Fried, Ronghang Hu, Volkan Cirik, Anna Rohrbach, Jacob Andreas, Louis-Philippe Morency, Taylor Berg-Kirkpatrick, Kate Saenko, Dan Klein, and Trevor Darrell. Speaker-Follower Models for Vision-and-Language Navigation. In *Proceedings of NeurIPS*, 2018. 2
- [24] Alex Graves, Santiago Fernandez, Faustino Gomez, and Juergen Schmidhuber. Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks. In *Proceedings of ICML*, 2006. 4
- [25] Mandy Guo, Zihang Dai, Denny Vrandečić, and Rami Al-Rfou. Wiki-40b: Multilingual language model dataset. In *LREC 2020*, 2020. 4
- [26] Udo Hahn, Katrin Tomanek, Elena Beisswanger, and Erik Faessler. A proposal for a configurable silver standard. In *Proceedings of the Fourth Linguistic Annotation Workshop*, pages 235–242, Uppsala, Sweden, July 2010. Association for Computational Linguistics. 2, 3
- [27] Harlan Hile, Radek Grzeszczuk, Alan Liu, Ramakrishna Vedantam, Jana Kosecka, and Gaetano Borriello. Landmark-based pedestrian navigation with enhanced spatial reasoning. In Hideyuki Tokuda, Michael Beigl, Adrian Friday, A. J. Bernheim Brush, and Yoshito Tobe, editors, *Pervasive Computing*, pages 59–76, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg. 2

- [28] Sepp Hochreiter and Jürgen Schmidhuber. Long Short-Term Memory. *Neural Computation*, 1997. 6
- [29] Yicong Hong, Cristian Rodriguez, Yuankai Qi, Qi Wu, and Stephen Gould. Language and visual entity relationship graph for agent navigation. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 7685–7696. Curran Associates, Inc., 2020. 2
- [30] Bernd Huber, Daniel McDuff, Chris Brockett, Michel Galley, and Bill Dolan. Emotional dialogue generation using image-grounded language models. In *CHI*, 2018. 3
- [31] Gabriel Ilharco, Vihan Jain, Alexander Ku, Eugene Ie, and Jason Baldridge. Effective and general evaluation for instruction conditioned navigation using dynamic time warping. *NeurIPS Visually Grounded Interaction and Language Workshop*, 2019. 7
- [32] Aashi Jain, Mandy Guo, Krishna Srinivasan, Ting Chen, Sneha Kudugunta, Chao Jia, Yinfei Yang, and Jason Baldridge. MURAL: Multimodal, multitask representations across languages. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3449–3463, Punta Cana, Dominican Republic, Nov. 2021. Association for Computational Linguistics. 2, 4, 6
- [33] Dan Klein and Christopher D. Manning. Accurate unlexicalized parsing. In *ACL*, 2003. 7
- [34] Alexander Klippel and Stephan Winter. Structural salience of landmarks for route directions. In Anthony G. Cohn and David M. Mark, editors, *Spatial Information Theory*, 2005. 5
- [35] Noriyuki Kojima, Alane Suhr, and Yoav Artzi. Continual learning for grounded instruction generation by observing human following behavior. *Transactions of the Association of Computational Linguistics*, 2021. 3
- [36] Alexander Ku, Peter Anderson, Roma Patel, Eugene Ie, and Jason Baldridge. Room-Across-Room: Multilingual Vision-and-Language Navigation with Dense Spatiotemporal Grounding. In *Proceedings of EMNLP*, 2020. 2, 3, 6, 8
- [37] Alex Ku, Peter Anderson, Jordi Pont-Tuset, and Jason Baldridge. Pangea: The panoramic graph environment annotation toolkit. In *2nd Workshop on Advances in Language and Vision Research (ALVR) at NAACL*, 2021. 7
- [38] Taku Kudo and John Richardson. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2018. 5
- [39] Shuhei Kurita and Kyunghyun Cho. Generative language-grounded policy in vision-and-language navigation with bayes’ rule. In *ICLR*, 2021. 1
- [40] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *ECCV*, 2018. 5
- [41] T.Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollar, and C. L. Zitnick. Microsoft COCO: Common Objects in Context. In *ECCV*, 2014. 5
- [42] Vivien Mast and Diedrich Wolter. A probabilistic framework for object descriptions in indoor route instructions. In Thora Tenbrink, John Stell, Antony Galton, and Zena Wood, editors, *Spatial Information Theory*, pages 185–204, Cham, 2013. Springer International Publishing. 1
- [43] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hour-glass networks for human pose estimation. In *ECCV*, 2016. 5
- [44] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *ACL*, 2002. 7
- [45] Alexander Pashevich, Cordelia Schmid, and Chen Sun. Episodic Transformer for Vision-and-Language Navigation. In *ICCV*, 2021. 2
- [46] Jordi Pont-Tuset, Jasper Uijlings, Soravit Changpinyo, Radu Soricut, and Vittorio Ferrari. Connecting vision and language with localized narratives. In *ECCV*, 2020. 3
- [47] Matt Post. A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, Brussels, Belgium, 2018. Association for Computational Linguistics. 7
- [48] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020. 2, 5
- [49] Dietrich Rebholz-Schuhmann, Antonio José Jimeno Yepes, Erik M. van Mulligen, Ning Kang, Jan Kors, David Milward, Peter Corbett, Ekaterina Buyko, Katrin Tomanek, Elena Beisswanger, and Udo Hahn. The CALBC silver standard corpus for biomedical named entities — a study in harmonizing the contributions from four independent named entity taggers. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC’10)*, Valletta, Malta, May 2010. European Language Resources Association (ELRA). 2, 3
- [50] Kai-Florian Richter and Alexander Klippel. A model for context-specific route directions. In Christian Freksa, Markus Knauff, Bernd Krieg-Brückner, Bernhard Nebel, and Thomas Barkowsky, editors, *Spatial Cognition IV. Reasoning, Action, Interaction*, pages 58–78, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg. 1
- [51] Michael Roth and Anette Frank. Computing EM-based alignments of routes and route directions as a basis for natural language generation. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 958–966, Beijing, China, Aug. 2010. Coling 2010 Organizing Committee. 1
- [52] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany, Aug. 2016. Association for Computational Linguistics. 5
- [53] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*, 2018. 6
- [54] Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer,

- and Dieter Fox. ALFRED: A Benchmark for Interpreting Grounded Instructions for Everyday Tasks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [55] Michael Sollami and Aashish Jain. Multimodal conditionality for natural language generation, 2021. 3
- [56] Lucia Specia, Stella Frank, Khalil Sima'an, and Desmond Elliott. A shared task on multimodal machine translation and crosslingual image description. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 543–553, Berlin, Germany, Aug. 2016. Association for Computational Linguistics. 6
- [57] Alane Suhr, Claudia Yan, Jack Schluger, Stanley Yu, Hadi Khader, Marwa Mouallem, Iris Zhang, and Yoav Artzi. Executing instructions in situated collaborative interactions. 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 2119–2130, Hong Kong, China, Nov. 2019. Association for Computational Linguistics. 3
- [58] Hao Tan, Licheng Yu, and Mohit Bansal. Learning to navigate unseen environments: Back translation with environmental dropout. In *NAACL*, 2019. 1, 2, 6, 7
- [59] Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *CVPR*, 2015. 7
- [60] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision, 2021. 3
- [61] Qingrong Xia, Zhenghua Li, Rui Wang, and Min Zhang. Stacked AMR parsing with silver data. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4729–4738, Punta Cana, Dominican Republic, Nov. 2021. Association for Computational Linguistics. 2, 3
- [62] Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. mT5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online, June 2021. Association for Computational Linguistics. 2, 5
- [63] Demet Yesiltepe, Ruth Dalton, and Ayse Torun. Landmarks in wayfinding: a review of the existing literature. In *Cognitive Processing*, 2021. 2
- [64] Ming Zhao, Peter Anderson, Vihan Jain, Su Wang, Alex Ku, Jason Baldridge, and Eugene Ie. On the evaluation of vision-and-language navigation instructions. In *EACL*, 2021. 2, 7
- [65] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as Points. In *Proceedings of CVPR*, 2019. 5
- [66] Yuanen Zhou, Meng Wang, Daqing Liu, Zhenzhen Hu, and Hanwang Zhang. More grounded image captioning by distilling image-text matching model. In *CVPR*, 2021. 3