

An Empirical Study on the Generalization Power of Neural Representations Learned via Visual Guessing Games

Alessandro Suglia¹, Yonatan Bisk², Ioannis Konstas¹, Antonio Vergari³,
Emanuele Bastianelli¹, Andrea Vanzo¹, and Oliver Lemon¹

¹Heriot-Watt University, Edinburgh, UK

²Carnegie Mellon University, Pittsburgh, USA

³University of California, Los Angeles, USA

¹{as247, i.konstas, e.bastianelli, a.vanzo, o.lemon}@hw.ac.uk

²ybisk@cs.cmu.edu, ³aver@cs.ucla.edu

Abstract

Guessing games are a prototypical instance of the “learning by interacting” paradigm. This work investigates how well an artificial agent can benefit from playing guessing games when later asked to perform on novel NLP downstream tasks such as Visual Question Answering (VQA). We propose two ways to exploit playing guessing games: 1) a supervised learning scenario in which the agent learns to mimic successful guessing games and 2) a novel way for an agent to play by itself, called Self-play via Iterated Experience Learning (SPIEL). We evaluate the ability of both procedures to generalise: an in-domain evaluation shows an increased accuracy (+7.79) compared with competitors on the evaluation suite CompGuessWhat?!; a transfer evaluation shows improved performance for VQA on the TDIUC dataset in terms of harmonic average accuracy (+5.31) thanks to more fine-grained object representations learned via SPIEL.

1 Background & Related Work

Learning a language requires interacting with both the environment and other agents (Bisk et al., 2020). Language games represent one common example of this (Wittgenstein et al., 1953), as seen by the important role of play in L1 child language acquisition (Hainey et al., 2016) as well as L2 learners (Godwin-Jones, 2014).

Among the language games defined in the literature (Steels, 2015), *guessing games* represent the first step in a curriculum for language learning. For example, in GuessWhat?! (de Vries et al., 2017), two agents interact with each other: a *Questioner* generates questions aimed at finding a hidden object in the scene and an *Oracle*, aware of the target object, answers the questions supporting the Questioner in playing the game. Different from other language games (Das et al., 2017), guessing games

have a specific goal which represents a clear incentive for learning. In addition, they require that the *Questioner* masters both natural language generation and understanding with a focus on object categories and attributes. For humans, concepts learned in this way are generic and generalisable to new tasks and domains where grounded reasoning is important (Hampton, 1979). However, *how well can AI agents generalise with concepts acquired from visual guessing games?*

The literature has not explored if representations built from self-play are transferable, focusing instead on large scale self-supervised learning. For instance, large scale image captioning datasets have been used to train multi-modal Transformers (Lu et al., 2019; Li et al., 2019; Tan and Bansal, 2019; Chen et al., 2019). Multi-task learning (Lu et al., 2020) has been used to leverage the diversity of training signals provided combining datasets, but only for discriminative tasks. While some dialogue work (Cogswell et al., 2020) aims to bootstrap a conversing agent from VQA datasets, most work on GuessWhat?! (de Vries et al., 2017; Shekhar et al., 2019; Strub et al., 2017) has designed bespoke models for the task, ignoring the utility of this dataset for other Vision+Language tasks.

We propose self-play as a mechanism for learning general grounded representations. We seed our approach with the *GuessWhat?!* corpus of questions and objects, and demonstrate how to generalise to other downstream tasks. We propose two different strategies to exploit these data. First, a supervised learning phase is undertaken to learn a Questioner and Oracle model able to play guessing games. Second, the trained agents can be used to play guessing games on images requiring only object annotations as supervision. We show that an agent trained on GuessWhat?! dialogues can use self-play to adapt to new and harder tasks. Specifically, we investigate models’ gener-

alisation performance and quality of the learned representations on the CompGuessWhat?! benchmark (Suglia et al., 2020), a more extensive evaluation suite for GuessWhat?!. Furthermore, we study how the learned representation help solve VQA on the dataset TDIUC (Kafle and Kanan, 2017). We show overall comparable performance with state-of-the-art models and improvements for specific question types that require object attribute information to be answered correctly.

2 Methodology

Our proposed transfer/fine-tuning procedure requires a training set of guessing games D_g from which we learn a Questioner Q and an Oracle O via supervised learning. Given a set of images \mathcal{I} , it is possible to use the trained models Q and O to run the self-play procedure for n epochs obtaining the model Q^n . Finally, given a downstream task t and an associated dataset \mathcal{D}_t based on images from \mathcal{I} , we use Q^n 's parameters as initialisation for the training procedure on \mathcal{D}_t .

To apply this procedure, both the *Questioner* and the *Oracle* require a *multi-modal encoder* Γ able to generate d -dimensional representations for the textual tokens \mathbf{h}_t , for the objects \mathbf{h}_o , as well as fusing the visual and textual modalities in a representation of the current context \mathbf{h}_c . After the self-play procedure, only the encoder Γ of the model Q^n is used in the fine-tuning process on the downstream task t using the dataset \mathcal{D}_t . It is important to underline that the presented self-play procedure does not depend on a specific implementation of the multi-modal encoder Γ . A possible implementation is presented in Section 2.4 and it is used in the experimental evaluation of this paper.

2.1 Oracle design

The Oracle task is cast as a Visual Question Answering (VQA) task conditioned on the image I , the current question q and on the target object \hat{o} . We follow common practice in vocabulary-based VQA (Antol et al., 2015) and we treat the problem as a multi-class classification task over the classes $\{Yes, No, N/A\}$. We use \mathbf{h}_c as input to a multi-layer feedforward neural network to obtain a probability distribution over the label set.

2.2 Questioner design

The Questioner must play two roles: question generation and target object prediction (de Vries et al.,

2017). It is beneficial to jointly learn the two tasks because the representations learned by each task are complementary. In addition, they better encode attributes, which favours better generalisation to unseen object categories (Suglia et al., 2020).

To solve the two specific tasks in a multi-task fashion, we design two different heads on top of the shared encoder Γ : 1) the *guesser head*, produces a probability distribution over every object o_i using the encoded representations \mathbf{h}_{o_i} passed through an MLP; 2) the *generator head*, a multi-modal decoder, also implemented as an MLP, which predicts a probability distribution over the vocabulary V given the context representation generated by Γ .

We include two losses in our model: 1) the negative log-likelihood of the probability associated by the guesser head with the target object \hat{o} (Shekhar et al., 2019); 2) a sequence-to-sequence cross-entropy loss (Sutskever et al., 2014) for the generated question tokens. Unlike previous work that trains a separate module to learn to stop (Shekhar et al., 2018), we add a special token [STOP] to the input data so that it learns when to stop more efficiently as part of the question generation task.

Training an agent to solve tasks of different complexity and size is challenging. The procedure presented in (Shekhar et al., 2019) alternates between tasks, updating the hardest task more often. For this technique, finding the right schedule is cumbersome and requires fine-tuning. We rely on a more systematic training procedure based on random dataset-proportional batch sampling inspired by (Sanh et al., 2019). This represents a hard-parameter sharing multi-task training procedure that avoids interference between tasks and favours a more stable training, which mitigates catastrophic forgetting (French, 1999).

2.3 Self-Play via Iterated Experience Learning (SPIEL)

Inspired by iterated learning (Kirby et al., 2014), we design a process by which the Questioner learns from games previously generated by other instances of the Questioner agent. We call our training procedure *Self-play via Iterated Experience Learning* (SPIEL).

In SPIEL, described in Algorithm 1, we assume access to a set of images \mathcal{I} and the bounding boxes \mathcal{O}_I of the objects therein.¹ In every gameplay, there

¹Object annotations intended as either gold bounding boxes or predicted bounding boxes from an object detector.




Figure 1: We use the single-stream VLP model as a backbone multi-modal encoder for our task. The visual features tokens (marked in red) are the FastRCNN features associated with the objects in the image, the history tokens (marked in blue) and the tokens to be generated (marked in yellow) are given in input to the model. A *Guesser head* uses the learned contextual object representations to generate a probability distribution over the objects $P(o_i|h_{o_i})$, whereas the *Generator head* is used to incrementally predict the masked tokens.

Algorithm 1 SPIEL: Self-Play via Iterated Experience Learning

```

1: procedure SELF_PLAY( $Q_0, O, \mathcal{I}, n$ )
2:    $\mathcal{D}_q \leftarrow$  READ_GOLD_GAMES()
3:    $\mathcal{E}_g \leftarrow \square$   $\triangleright$  Initialise the experience buffer
4:   for  $e \leftarrow 1, n$  do
5:      $\triangleright$  Interactive phase
6:      $Q \leftarrow Q^e$   $\triangleright$  load latest weights
7:      $\mathcal{G}^e \leftarrow$  GENERATE_GAMES( $\mathcal{I}$ )
8:      $\mathcal{G}^e \leftarrow$  PLAY_GAMES( $Q, O, \mathcal{G}^e$ )
9:     APPEND( $\mathcal{E}_g, \mathcal{G}^e$ )
10:     $\mathcal{D}_g^e \leftarrow \square$ 
11:     $\triangleright$  Transmission phase
12:    for  $i \leftarrow 0, \text{len}(\mathcal{E}_g)$  do
13:       $g \leftarrow \mathcal{E}_g[i]$   $\triangleright$  Priority to the latest games
14:      if IS_VALID_GAME( $g$ ) then
15:        APPEND( $\mathcal{D}_g^e, g$ )
16:      if LEN( $\mathcal{D}_g^e$ ) == LEN( $\mathcal{D}_q$ ) then break
17:     $\triangleright$  Learning phase
18:     $Q^{e+1} \leftarrow$  TRAIN( $Q, \mathcal{D}_q, \mathcal{D}_g^e$ )

```

is a Questioner Q and an Oracle O , initialised with agents Q^0 and O , respectively, that were trained with Supervised Learning using gold successful dialogues.² We consider every iteration e of the algorithm as a *self-play* epoch. In a single self-play epoch, we alternate 3 phases:

Interactive phase: the agents play guessing games with novel combinations of image and target object. The generated dialogue can be successful if the predicted target object is equal to the target object. Every played dialogue is stored in an experience buffer \mathcal{E}_g .

Transmission phase: in this phase the datasets for the multi-task learning procedure for the Questioner are created. The generator head dataset \mathcal{D}_q is fixed in advance while the dataset for the guesser head \mathcal{D}_g^e is created from the experience buffer \mathcal{E}_g by selecting the *unique* and *valid* dialogues.

²The Oracle is fixed during this learning procedure.

Learning phase: the same multi-task learning procedure used in the supervised learning phase is used to fine-tune the Questioner parameters using the datasets \mathcal{D}_g^e and \mathcal{D}_q collected for the current epoch e . This procedure is repeated n times or until a halting condition is reached (e.g. early stopping based on validation metric).

See Appendix A.1 for implementation details. At the end of the SPIEL procedure, we obtain the model Q^n whose parameters can be reused in other tasks. Particularly, we use the parameters of Q^n 's shared encoder Γ as initialisation for the fine-tuning on the downstream task t using dataset \mathcal{D}_t .

2.4 Implementation

We implement a shared multi-modal encoder Γ using VLP (Zhou et al., 2020), a single-stream multi-modal Transformer for captioning depicted in Figure 1. During the GuessWhat?! fine-tuning, we extend VLP by including dialogue context in the input together with the features associated with the objects in the image. We learn two new segment ids to represent the question/answer exchanges in the dialogue, as described in (Wolf et al., 2019). The question is generated by incrementally replacing [MASK] tokens until the end of sequence is generated. See Appendix A.2 for more details. SPIEL training is run on a set of images \mathcal{I} from GuessWhat?! and TDIUC dataset with corresponding object annotations. We make sure that GuessWhat?! test images are not contained in \mathcal{I} . This is not an issue for TDIUC test images because the downstream task annotations (QA pairs) are not used by the model during this phase. Once the model has been trained with SPIEL, we use the parameters of the shared encoder Γ as a backbone for a VQA model that is fine-tuned on the TDIUC dataset.

3 Experimental Evaluation

To assess the generality of our learned representations, we include two evaluation paradigms: 1) *in-domain evaluation* and 2) *transfer evaluation*. We evaluate several variants of our model: 1) VLP-SL: VLP-based model trained on GuessWhat?! data using multi-task learning; 2) SPIEL-gs: VLP-SL model fine-tuned with our SPIEL procedure where the generator head uses only gold successful games (gs); 3) SPIEL-gm: same as 2) but both successful and failed gold games are used by the generator head. In both SPIEL variants, the guesser head is trained using failed and successful generated games because it is important for the guesser head to be exposed to both types of signal to learn a more robust policy. We decided to investigate the two variants SPIEL-gs and SPIEL-gm to get more insights about the effect that successful and failed games have on the generator head ability to produce effective dialogues.

3.1 In-domain evaluation

We use the CompGuessWhat?! evaluation suite (Suglia et al., 2020) to assess the ability of the Questioner to play guessing games and learn visually grounded representations in the process. It complements an evaluation based only on gameplay accuracy (de Vries et al., 2017) with 2 auxiliary tasks: target object 1) attribute-prediction expressed in terms of abstract attributes (A), situated-attributes (SO), abstract+situated attributes (AS), and location attributes (L); 2) zero-shot gameplay with near-domain accuracy (ND) and out-of-domain accuracy (OD). Table 1 shows the comparison with previous state-of-the-art models on this benchmark such as de Vries et al. (2017) (DV-*) and Shekhar et al. (2019) (GDSE-*). VLP-SL has a greater advantage in terms of representation power compared to previous models. This is reflected in all the tasks of the CompGuessWhat?! evaluation. Particularly, we see better performance even for the zero-shot gameplay (ND: +5.6, OD: +15.2). This is because VLP associates a vector of probabilities that represents a distribution over the VisualGenome object classes with every object. This helps VLP to cope with the issue of unseen objects and helps the model to generalise. Learning to play is key to gameplay performance, leading to an increase of +4.4 over VLP-SL and +7.9 over GDSE-CL. In this setup, the difference between the versions SPIEL-gs and SPIEL-gm is very

Models	Acc.	Attribute Pred.				ZShot		Score
		A	SO	AS	L	ND	OD	
Random	15.8	15.1	0.1	7.8	2.8	16.8	18.6	13.3
DV-SL	41.5	46.8	39.1	48.5	42.7	31.3	28.4	38.5
DV-RL	53.5	45.2	38.9	47.2	43.5	43.9	38.7	46.2
GDSE-SL	49.1	59.9	47.6	60.1	48.3	29.8	22.3	43.0
GDSE-CL	59.8	59.5	47.6	59.8	48.1	43.4	29.8	50.1
VLP-SL	59.5	59.2	48.2	59.7	49.3	49.0	45.0	53.5
SPIEL-gs	64.1	61.3	49.6	61.6	51.1	54.9	51.9	57.8
SPIEL-gm	64.6	60.8	48.3	59.5	51.0	55.3	52.9	57.9

Table 1: F1 scores for attribute prediction and accuracies for zero-shot evaluation on CompGuessWhat?!.

minimal (0.1). However, when analysed in more detail, we can see that training the questioner with gold successful data only improves attribute prediction while using mixed data improves overall generalisation in the zero-shot evaluation.

3.2 Transfer evaluation

For the transfer evaluation, we use the VQA dataset TDIUC (Kafle and Kanan, 2017). It provides a finer-grained way to assess the quality of the representations learned by our guessing game transfer technique in terms of several question types including object categories and their attributes. Specifically, we were interested in improving on the following question types: 1) Positional reasoning; 2) Counting; 3) Object presence; 4) Utility/Affordances; 5) Attribute; 6) Color; and 7) Object recognition. TDIUC is evaluated using the arithmetic mean accuracy per question type (A-MPT), as well as the harmonic mean (H-MPT) that better captures the skewed question-type distribution. In Table 2, we report a comparison between variants trained on guessing games data (VLP+SL and SPIEL-*), the original model VLP trained on Conceptual Captions (VLP+CC) and other state-of-the-art models *specifically* designed for the VQA task such as MUREL (Cadene et al., 2019), RAU (Noh and Han, 2016), NMN (Andreas et al., 2016), MCB-* (Fukui et al., 2016). The full set of results is available in the Appendix, Table 4.

Among them, MUREL achieves the best scores across the board, due to a custom iterative reasoning mechanism and a non-linear fusion module. However, all our models have a more balanced overall performance which results in better harmonic means (H-MPT, +5 points over MUREL). Specifically, this improvement is favoured by an increase in accuracy on the *Utility/Affordances* ques-




Figure 2: We show the ability of the model to play guessing games with the *bowl* as target object (highlighted in red). Given the generated dialogue, we use the probing classifier trained for CompGuessWhat?! to predict the *bowl*'s attributes. Predictions on TDIUC questions associated with the current image are reported as well.

tion type (+20.7). As shown by the attribute prediction in the CompGuessWhat?! and depicted in Figure 2 (c), our models learn better representations than competitors specifically for abstract attributes among which there are object affordances. Particularly, we can see how it is able to understand that certain objects can *contain* things (e.g. “the one with the soup in it?”), that objects have specific functions (e.g. “are the contents of the plate edible?”) or that they have specific properties (e.g. “a spoon is made of wood”).

The effectiveness of the proposed fine-tuning procedure is confirmed by the improved performance across all the question types compared to our baseline VLP+CC. Models such as MUREL and MCB- \star equipped with specific VQA modules have an advantage on specific question (e.g., positional reasoning) compared to VLP that relies only on BERT self-attention layers (Devlin et al., 2019). In addition, when comparing the two SPIEL variants, a similar trend showed in the in-domain evaluation can be observed. Particularly, SPIEL-gm benefits from being exposed to more language data coming from successful and failed guessing games.

4 Conclusions

In this work, we verified that representations learned while playing guessing games can be transferred to other downstream tasks such as VQA. We presented two ways of learning from guessing games data namely multi-task learning and SPIEL. Models using SPIEL performed better both on in-domain evaluation on CompGuessWhat?! as well as on the transfer task TDIUC. Our self-play pro-

MODEL	POSITION	COUNT	PRESENCE	AFFORD.	ATTR.	COLOR	RECOG.	A-MPT	H-MPT
RAU	35.3	48.4	94.4	31.6	56.5	66.9	86.1	67.8	59.0
NMN	27.9	49.2	92.5	25.2	47.7	54.9	82.0	62.6	51.9
MCB-A	55.4	51.0	93.6	35.1	56.7	68.5	81.9	67.9	60.5
MCB	33.3	50.3	91.8	33.9	53.2	56.9	84.6	65.8	58.0
MUREL	41.2	61.8	95.8	21.4	58.2	74.4	89.4	71.2	59.3
VLP									
+CC	36.9	55.3	94.7	31.0	55.4	67.3	85.8	68.8	60.1
+SL	39.0	57.6	94.8	42.1	54.3	69.0	86.1	70.5	64.0
SPIEL-gs	40.9	57.5	94.8	36.3	56.9	69.2	86.3	70.4	63.3
SPIEL-gm	40.6	57.0	94.8	39.2	57.0	69.4	86.2	70.9	64.3

Table 2: Results for the transfer evaluation on TDIUC. The models are divided in two categories: (top) Models specifically designed for VQA and (bottom) our VLP-based implementations. We report only the question types that we believe will benefit from the guessing games fine-tuning procedure. For the full set of results please refer to Appendix, Table 4.

cedure was able to learn useful and finer-grained object representations such as object affordances, thus demonstrating that *learning to guess helps learning to ground*.

The current study showed how we can apply the SPIEL training procedure to a VQA dataset such as TDIUC. We believe that this work can be extended to other datasets because the SPIEL procedure only requires a set of images and associated object bounding boxes. These could be either gold or generated by a trained object detector therefore classifying guessing games as a *holistic self-training procedure* for multi-modal datasets.

References

- Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. 2016. Neural module networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 39–48.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433.
- Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, et al. 2020. Experience grounds language. *arXiv preprint arXiv:2004.10151*.
- Remi Cadene, Hedi Ben-Younes, Matthieu Cord, and Nicolas Thome. 2019. Murel: Multimodal relational reasoning for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1989–1998.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2019. Uniter: Learning universal image-text representations. *arXiv preprint arXiv:1909.11740*.
- Michael Cogswell, Jiasen Lu, Rishabh Jain, Stefan Lee, Devi Parikh, and Dhruv Batra. 2020. Dialog without dialog data: Learning visual dialog agents from vqa data. *arXiv preprint arXiv:2007.12750*.
- Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José MF Moura, Devi Parikh, and Dhruv Batra. 2017. Visual dialog. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 326–335.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. In *Advances in Neural Information Processing Systems*, pages 13063–13075.
- Robert M French. 1999. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135.
- Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. 2016. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 457–468.
- Robert Godwin-Jones. 2014. Games in language learning: Opportunities and challenges. *Language Learning & Technology*, 18(2):9–19.
- Thomas Hainey, Thomas M Connolly, Elizabeth A Boyle, Amanda Wilson, and Aisya Razak. 2016. A systematic literature review of games-based learning empirical evidence in primary education. *Computers & Education*, 102:202–223.
- James A Hampton. 1979. Polymorphous concepts in semantic memory. *Journal of verbal learning and verbal behavior*, 18(4):441–461.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*.
- Kushal Kafle and Christopher Kanan. 2017. An analysis of visual question answering algorithms. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1965–1973.
- Simon Kirby, Tom Griffiths, and Kenny Smith. 2014. Iterated learning and the evolution of language. *Current opinion in neurobiology*, 28:108–114.
- Jason Lee, Kyunghyun Cho, and Douwe Kiela. 2019. Countering language drift via visual grounding. 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 4376–4386.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems*, pages 13–23.
- Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 2020. 12-in-1: Multi-task vision and language representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10437–10446.
- Hyeonwoo Noh and Bohyung Han. 2016. Training recurrent answering units with joint loss minimization for vqa. *arXiv preprint arXiv:1606.03647*.
- Victor Sanh, Thomas Wolf, and Sebastian Ruder. 2019. A hierarchical multi-task approach for learning embeddings from semantic tasks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6949–6956.

- Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. 2015. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565.
- Ravi Shekhar, Tim Baumgärtner, Aashish Venkatesh, Elia Bruni, Raffaella Bernardi, and Raquel Fernández. 2018. Ask no more: Deciding when to guess in referential visual dialogue. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1218–1233.
- Ravi Shekhar, Aashish Venkatesh, Tim Baumgärtner, Elia Bruni, Barbara Plank, Raffaella Bernardi, and Raquel Fernández. 2019. Beyond task success: A closer look at jointly learning to see, ask, and guess-what. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2578–2587.
- Luc Steels. 2015. *The Talking Heads experiment: Origins of words and meanings*, volume 1. Language Science Press.
- Florian Strub, Harm De Vries, Jeremie Mary, Bilal Piot, Aaron Courville, and Olivier Pietquin. 2017. End-to-end optimization of goal-driven and visually grounded dialogue systems. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 2765–2771.
- Alessandro Suglia, Ioannis Konstas, Andrea Vanzo, Emanuele Bastianelli, Desmond Elliott, Stella Frank, and Oliver Lemon. 2020. [CompGuess-What?!](#): A multi-task evaluation framework for grounded language learning. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7625–7641, Online. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Hao Tan and Mohit Bansal. 2019. Lxmert: Learning cross-modality encoder representations from transformers. 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 5103–5114.
- Harm de Vries, Florian Strub, Sarath Chandar, Olivier Pietquin, Hugo Larochelle, and Aaron C. Courville. 2017. [Guesswhat?! visual object discovery through multi-modal dialogue](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 4466–4475. IEEE Computer Society.
- Ludwig Wittgenstein, Gertrude Elizabeth Margaret Anscombe, and Rush Rhees. 1953. *Philosophische Untersuchungen. (Philosophical investigations)*.
- Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. Transfertransfo: A transfer learning approach for neural network based conversational agents. *arXiv preprint arXiv:1901.08149*.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason J Corso, and Jianfeng Gao. 2020. Unified vision-language pre-training for image captioning and vqa. In *AAAI*, pages 13041–13049.
- Bohan Zhuang, Qi Wu, Chunhua Shen, Ian Reid, and Anton Van Den Hengel. 2018. Parallel attention: A unified framework for visual object discovery through dialogs and queries. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4252–4261.

A Appendices

A.1 Self-Play via Iterated Experience Learning (SPIEL)

Learning to replicate gold dialogues is not enough to play successfully. High performance in gameplay can be achieved only when the agents start playing the game and are exposed to their own mistakes. Reinforcement Learning (Strub et al., 2017) or Collaborative Learning (Shekhar et al., 2019) are possible approaches to tackle this problem.

Inspired by iterated learning (Kirby et al., 2014), we design a process by which “the gameplay arises in one instance of the questioner through induction on the basis of observations of gameplay in other questioner agents who acquired that gameplay capability in the same way”. Therefore, we call our procedure *Self-play via Iterated Experience Learning* (SPIEL).

In this setup, we assume we have access to a set of images \mathcal{I} and for each image I we have object bounding boxes \mathcal{O}_I . The SP training procedure, showed in Figure 1, can be described as follows. We assume that there is a Questioner agent Q and

an Oracle agent O . At the beginning of the procedure they are initialised with agents Q^0 and O , respectively, trained with Supervised Learning using gold successful dialogues³. We consider every iteration e of the algorithm as a *self-play* epoch. In a single self-play epoch we alternate 3 phases: 1) *interactive* phase: the agents play guessing games with novel combinations of image and target object; 2) *transmission* phase: the questioner creates new datasets from the dialogues generated over the epochs; 3) *learning* phase: multi-task learning is used to fine-tune the Questioner parameters using the datasets collected for the current epoch.

A.1.1 Interactive phase

We start the interactive phase by first sampling a set of reference games \mathcal{G}^e which consists of pairs (I, \hat{o}) where $I \in \mathcal{I}$ and \hat{o} is the target object sampled at random from the object annotations \mathcal{O}_I . The agents Q^e and O play the games \mathcal{G}^e and accumulate the generated experiences. During this phase, the questioner agent is using the most updated weights generated at epoch $e - 1$. It generates questions by nucleus sampling (Holtzman et al., 2019) from the probability distribution over the vocabulary learned by the generator head. When the [STOP] token is sampled, the guesser head, conditioned on the dialogue generated so far, selects the object \tilde{o} with the highest probability. A game is *successful* if the predicted object \tilde{o} is equal to the target object \hat{o} .

A.1.2 Transmission phase

For every epoch e , in the transmission phase, we create the datasets \mathcal{D}_q and \mathcal{D}_g for the questioner and guesser heads, respectively, used in the learning phase for the questioner parameters update.

Questioner experience buffer To make sure that the questioner does not experience *language drift* (Lee et al., 2019), we consider a fixed dataset \mathcal{D}_q composed of dialogues generated by humans contained in the GuessWhat?! training data. The shared encoder Γ benefits from this data too because it is still exposed to human generated language, which guarantees better generalisation.

Guesser experience buffer The Guesser should learn from its own mistakes – therefore we use generated dialogues for the model updates (de Vries et al., 2017; Shekhar et al., 2019). Inspired by Prioritised Experience Replay (Schaul et al., 2015), we create the experience buffer for the guesser \mathcal{E}_g^e by accumulating all the *unique* and *valid* dialogues

³The Oracle is fixed during this learning procedure.

generated until epoch e . We consider a dialogue *unique* if \mathcal{D}_g^e does not contain another dialogue with the same encoding⁴. In addition, we consider a dialogue *valid* if it does not contain repeated questions. We cap the number of dialogues in \mathcal{D}_g^e so that it matches the number of experiences in \mathcal{D}_q . This is done so that during the multi-task training procedure there is an equal number of dialogues for each task from which the agent will learn.

A.1.3 Learning phase

In this phase, we use the same multi-task training procedure that was used during the supervised learning phase. We update the Questioner parameters using the dialogues collected in \mathcal{D}_q and \mathcal{D}_g^e . The updated parameters resulting from this step will be used for the self-play epoch $e + 1$.

A.2 VLP implementation

A.2.1 Multi-modal encoder

To implement the agents in our guessing games, we rely on VLP, a single-stream multi-modal model (Zhou et al., 2020) that jointly learns visual and language representations using Conceptual Captions (CC) dataset (Sharma et al., 2018). The input starts with a classification token ([CLS]), followed by a series of K visual tokens, a separation token ([SEP]) divides the *dialogue* sequence from the visual and from the sequence of *tokens to be generated*. In a guessing game, we represent the reference image I as a set of image regions extracted from an off-the-shelf object detector $\{r_1, r_2, \dots, r_K\}$. Following (Zhou et al., 2020), each region r_i is represented by linear transformation of a feature vector $f \in \mathbb{R}^{d_n}$, region class probabilities $c \in \mathbb{R}^{d_c}$ and region geometric information $g \in \mathbb{R}^{d_o}$ where $d_o = 5$ consists of four values for top left and bottom right corner coordinates of the region bounding box (normalized between 0 and 1) and one value for its relative area (i.e., ratio of the bounding box area to the image area, also between 0 and 1). The Questioner models uses at most 36 predicted bounding boxes from FastRCNN while the Guesser is using features generated by FastRCNN for gold bounding boxes. We use a specific segment id s_v for every region.

For the language part, we use Wordpiece embeddings (Wu et al., 2016). In particular, we flatten the turns of the dialogue context as a sequence of

⁴The encoding of a dialogue is the SHA-256 hash associated with its sequence of tokens.

tokens. However, to allow the model to differentiate between question and answer tokens, following (Wolf et al., 2019), we rely on novel segment ids (s_u, s_a). The VLP’s hidden state of the [CLS] token is used as context representation \mathbf{h}_c .

A.2.2 Oracle design

The implementation of the Oracle follows the one presented in the original VLP paper to solve the VQA task (Zhou et al., 2020). Particularly, the model predicts a probability distribution over the possible answers by using a multi-layer feed-forward neural network that receives in input the element-wise product between the hidden state associated with the [CLS] token and the hidden state associated with target object. The model is optimised by minimising the cross-entropy loss using as training dataset the question/answer pairs in the successful GuessWhat?! training dialogues.

A.2.3 Questioner design

We rely on the VLP ability to generate captions for the question generation task. In particular, we provide in input to the model: 1) predicted FastRCNN visual features following (Zhou et al., 2020); 2) dialogue generated so far as a flattened sequence of tokens; 3) question to be generated. We use another segment id s_q to allow the model to differentiate what is the input and which are the tokens to be generated. Following (Dong et al., 2019), we make sure that the attention mask for tokens of the question to be generated are masked so that the token at timestep t is not allowed to attend to the future tokens (*seq2seq attention mask*). For this specific model, we use the masked language modelling objective (Devlin et al., 2019) casting the task as multi-modal masked language modelling.

A.3 GuessWhat?! evaluation

Oracle evaluation We report the test accuracy for the Oracle of 82.22%. The baseline model used by all the other is 78.5% (de Vries et al., 2017).

Guesser evaluation We report in Table 3 the accuracy of the guesser in predicting the target object when gold dialogues are given in input. We compare this model with several baselines reported in (de Vries et al., 2017) (first block), more sophisticated methods such as `ParallelAttention` (Zhuang et al., 2018) and `GDSE-*` (Shekhar et al., 2019) (second block) as well as other Transformer-based models such as `VILBERT` (Lu et al., 2020) (third block).

Model	Accuracy
Human	90.80%
Random	17.10%
LSTM	61.30%
HRED	61%
LSTM+VGG	60.50%
HRED+VGG	60.40%
ParallelAttention	63.40%
GDSE-SL	62.96%
GDSE-CL	59.79%
VILBERT	65.69%
VLP-SL	69.30%
SPIEL-gs	71.80%
SPIEL-gm	71.70%

Table 3: Results for the guesser accuracy evaluation on gold dialogues.

MODEL	<u>POSITIONAL</u>	<u>COUNTING</u>	<u>PRESENCE</u>	<u>AFFORDANCES</u>	<u>ATTRIBUTE</u>	<u>COLOR</u>	<u>RECOGNITION</u>	SCENE	ABSURD	SENTIMENT	ACTIVITY	SPORT	ACCURACY	A-MPT	H-MPT
MUREL	41.19	61.78	95.75	21.43	58.19	74.43	89.41	96.11	99.80	60.65	63.83	96.20	88.20	71.20	59.30
RAU	35.26	48.43	94.38	31.58	56.49	66.86	86.11	93.96	96.08	60.09	51.60	93.47	84.26	67.81	59.00
NMN	27.92	49.21	92.50	25.15	47.66	54.91	82.02	91.88	87.51	58.02	44.26	89.99	79.56	62.59	51.87
MCB-A	55.40	51.01	93.64	35.09	56.72	68.54	85.54	93.06	84.82	66.25	52.35	92.77	81.86	67.90	60.47
MCB	33.34	50.29	91.84	33.92	53.24	56.93	84.63	92.04	83.44	65.46	51.42	92.47	79.20	65.75	58.03
VLP-CC	36.93	55.28	94.65	30.99	55.42	67.33	85.76	92.98	98.34	62.62	51.34	94.11	85.60	68.81	60.14
VLP-SL	39.04	57.61	94.79	42.11	54.29	69.01	86.07	93.39	97.54	65.77	52.39	94.34	85.98	70.53	63.95
SPIEL-gs	40.94	57.53	94.76	36.26	56.87	69.2	86.33	93.97	97.48	62.3	54.44	94.62	86.1	70.39	63.34
SPIEL-gm	40.6	57.01	94.77	39.18	56.97	69.42	86.21	93.72	97.19	66.09	55.29	94.18	86	70.89	64.31

Table 4: Summary of results for the transfer evaluation on TDIUC. The models are divided in two categories: (1) Models which are specifically designed for VQA (top) and (2) models that rely on the VLP encoder to generalise to different downstream tasks (bottom). We underline the question types that we believe will benefit from the guessing games transfer/fine-tuning procedure.