Meeting Readers' Expectations: Presenting Context and Content Effectively

Réka Mihálka and Simon Milligan

Exercise 1. Compare the two paragraphs below and decide which one has better flow.

VERSION 1. Human pose recovery has garnered a lot of research interest in the vision community. A wide range of applications such as augmented reality, virtual shopping, human–robot interaction, etc. use it. The availability of 3D pose supervision from large-scale datasets [19, 42, 57] has made recent advances in human pose recovery possible. Although in-studio benchmarks enable superior performance, the models usually suffer from poor cross-dataset performance. Domain bias [36, 27] is often induced by training on synthetic and in-studio datasets, which lack diversity in subject appearance, lighting, background variation, among others, thereby restricting generalizability. The question arises from this: across diverse data domains, how can we bridge performance gaps?

VERSION 2. Human pose recovery has garnered a lot of research interest in the vision community. It is extensively used in a wide range of applications such as augmented reality, virtual shopping, human–robot interaction, etc. Recent advances in human pose recovery is largely attributed to the availability of 3D pose supervision from large-scale datasets [19, 42, 57]. Although the models achieve superior performance on in-studio benchmarks, they usually suffer from poor cross-dataset performance. Synthetic and in-studio datasets lack diversity in subject appearance, lighting, background variation, among others, which is why training on these datasets can easily induce a domain bias [36, 27] and restrict generalizability. This poses the question: how can we bridge performance gaps across diverse data domains?

Adapted from: Ramesha Rakesh Mugaludi, Jogendra Nath Kundu, Varun Jampani, Venkatesh Babu R. "Aligning Silhouette Topology for Self-Adaptive 3D Human Pose Recovery" Advances in Neural Information Processing Systems 34 (NeurIPS 2021)

Exercise 2. Which structure is used in the paragraph below?

Deep Neural Networks (DNNs) have achieved unprecedented success in a wide range of applications due to their remarkably high accuracy [15]. However, this high performance stems from significant growth in DNN model size; i.e., massive overparameterization. Furthermore, these highly overparameterized models are known to be susceptible to the out-of-distribution (OOD) shifts encountered during their deployment in the wild [5]. This resource-inefficiency and OOD brittleness of state-of-the-art (SOTA) DNNs severely limits the potential applications DL can make an impact on.

Source: James Diffenderfer, Brian Bartoldson, Shreya Chaganti, Jize Zhang, Bhavya Kailkhura. “A Winning Hand: Compressing Deep Networks Can Improve Out-of-Distribution Robustness” Advances in Neural Information Processing Systems 34 (NeurIPS 2021)

Exercise 3. Fill in the information.

The worst-case optimal algorithms, however, tend to be too conservative in actual practice. (1)_____ **[Demonstrative pronoun]** is because true worst-case environments are quite rare in real-world applications. Rather, (2) _____ **[Repetition]** may have structures that are convenient for the learner, and it is desirable that the algorithm takes advantage of such structures to improve performance. To exploit such (3) _____ **[Repetition]**, two main categories of approaches have been studied: adapting to (nearly) stochastic environments and developing data-dependent regret bounds.

Source: Shinji Ito “Hybrid Regret Bounds for Combinatorial Semi-Bandits and Adversarial Linear Bandits”. Advances in Neural Information Processing Systems 34 (NeurIPS 2021).

Exercise 4. Can you identify the function of each sentence in this paragraph?

While we have shown benefits of our method, it has limitations and interesting future work. First, matching using the homography matrix calculated by RANSAC might not be the best option for our hypergraph propagation. It does not consider local features’ contextual cues when applying matching, and it has the offline homography calculation overhead. Recent deep-learning-based matching techniques such as SuperGlue [29] show performance improvement in many tasks in terms of both speed and accuracy and can also be helpful for our hypergraph propagation. Second, to achieve a better result, both hypergraph propagation and community selection can be combined with existing query expansion, diffusion, and spatial verification methods. For example, it is possible to adapt region diffusion [10] to the hypergraph model or adapt community selection on ordinary diffusion methods.

Source: Guoyuan An, Yuchi Huo, Sung-eui Yoon. Hypergraph Propagation and Community Selection for Objects Retrieval, Advances in Neural Information Processing Systems 34 (NeurIPS 2021)

Exercise 5. Write a topic sentence for the paragraph below.

[…] The authors conjecture that the sampling based optimization procedure of UMAP prevents the minimization of the supposed loss function, thus not reproducing the high-dimensional similarities in embedding space. They substantiate this hypothesis by qualitatively estimating the relative size of attractive and repulsive forces. In addition, they implement a BarnesHut approximation to the loss function (6) and find that it yields a diverged embedding. We analyze UMAP’s sampling procedure in depth, compute UMAP’s true loss function in closed form and

contrast it against the supposed loss in Section 5. Based on this analytic effective loss function, we can further explain Böhm et al. [4]'s empirical finding that the specific high-dimensional similarities provide little gain over the binary weights of a shared kNN graph, see Section 6. Finally, our theoretical framework leads us to a new tentative explanation for UMAP's success in Section 7.

Source: Sebastian Damrich, Fred A. Hamprecht. "On UMAP's True Loss Function." Advances in Neural Information Processing Systems 34 (NeurIPS 2021).

Exercise 6. Food for thought

Which sections provide context to the content of which other sections? List a few pairs.

What adjustments do you need to make to an introduction if you don't have a background section?

To what other aspects of the scientific process can you apply the principle of «context before content»?

Contact:

Réka @ Manuscript Manufacture: manuscriptmanufacture.com

Simon @Acalanser: acalanser.ch

Please fill in this survey to help us tailor our offer to your needs.