

ADVERTIMENT. L'accés als continguts d'aquesta tesi queda condicionat a l'acceptació de les condicions d'ús establertes per la següent llicència Creative Commons:  <https://creativecommons.org/licenses/?lang=ca>

ADVERTENCIA. El acceso a los contenidos de esta tesis queda condicionado a la aceptación de las condiciones de uso establecidas por la siguiente licencia Creative Commons:  <https://creativecommons.org/licenses/?lang=es>

WARNING. The access to the contents of this doctoral thesis it is limited to the acceptance of the use conditions set by the following Creative Commons license:  <https://creativecommons.org/licenses/?lang=en>



**Universitat Autònoma
de Barcelona**

**Exploring the role of Text in Visual Question Answering
on Natural Scenes and Documents**

A dissertation submitted by Rubèn Pérez Tito at
Universitat Autònoma de Barcelona to fulfil the de-
gree of **Doctor of Philosophy**.

Bellaterra, September 27, 2023

Director | **Dr. Ernest Valveny Llobet**
Universitat Autònoma de Barcelona
Dept. Ciències de la Computació
Centre de Visió per Computador

Thesis
committee | **Dr. Josep Lladós Canet**
Universitat Autònoma de Barcelona
Dept. Ciències de la Computació
Centre de Visió per Computador

Dr. Andrew David Bagdanov
Università degli Studi di Firenze
Dept. Ingegneria dell'Informazione
Florence, Italy

Dr. Jon Almazán
V7Labs
Barcelona, Spain

International
evaluators | **Dr. Ron Litman**
Amazon AI
Computer Vision AI Labs
Tel Aviv, Israel

Dr. Anjan Dutta
University of Surrey
PAI and CVSSP
Surrey, United Kingdom



This document was typeset by the author using \LaTeX 2 ϵ .

The research described in this book was carried out at the Computer Vision Center, Universitat Autònoma de Barcelona.

This work is licensed under Creative Commons Attribution-ShareAlike 4.0 International (CC BY-SA 4.0) © ⓘ 2020 by Rubèn Pérez Tito. You are free to copy and redistribute the material in any medium or format as long as you attribute its author. If you alter, transform or build upon this work, you may distribute the resulting work only under the same, similar or compatible license.

ISBN 978-84-124793-5-5

Printed by Ediciones Gráficas Rey, S.L.

To my friends, family and fiancée...

Acknowledgments

A thesis can not be written without the people you have around. I have met so many great people across my PhD and the least I can do is to acknowledge their help.

Ernest, your patience and assistance during my PhD have been invaluable. You guided me towards the most intriguing and promising research topics and provided invaluable corrections to papers, reports, and finally, this thesis. Your support over the years has played a pivotal role in my successful completion of the PhD, and for that, I am profoundly grateful. Dimos, I've always regarded you as my unofficial second supervisor. Our collaboration, discussions, and the rebuttal of research reviews have significantly contributed to my personal and professional growth. I will forever remember your advice to "Attack the reviewers instead of defend yourself against them". Moreover, I have to acknowledge that it was your passion for this field and your exceptional teaching skills in the Pattern Recognition subject that lit my interest in pursuing Artificial Intelligence as a professional career.

To all my colleagues, but especially Ali and Andrés. From the outset of our master's program at UAB, you have been my constant companions. Ali, your willingness to help the others is astonishing and heartwarming, from assisting the entire class with the Security and Information assignments the first time we met to spend tons of hours reviewing bugs in the code with me. Throughout the entire PhD you have helped me a lot, providing very insightful discussions, talking about crazy ideas (*Agronet*), playing squash to release some stress or just spending some time talking about life. I cherish all the time we've spent together and look forward to our mountain adventure, as I've been insisting for the past seven years. Andrés, our debates on all kind of topics ranging from consciousness, health, sports and veganism among others, along with our shared work experiences, have enriched my life in many ways. Thank you both for your unwavering support. To the rest of my colleagues, with whom I have spent a lot of time and I heartfully appreciate, Sounak, Sanket, Sergi, Moha, Lei, Raul, Khanh, Pau, Adrià, Andrea, Alvin, Ayan, and all the students who have come at CVC for an internship hailing from India, Italy and all over the world. Thanks also to postdocs and professors from whom I've learned invaluable knowledge Lluís, Marçal, Josep, Alícia, Oriol. Furthermore, I would like to highlight the fostering, supportive and enriching work environment of the entire CVC, encompassing research, administrative personnel and IT team. It's been a privilege to be part of this wonderful community.

Despite all the support I received at CVC, I couldn't have completed the PhD without all the help and warmth from my personal life. Angi and Xavier, your warm welcome made me feel like part of the family from the very beginning. I've also learned a lot from you, living closer with nature and realizing that money may not be that important. Iaia Tere and tieta Marta, you've also embraced me as if I were another grandchild and nephew. I cherish all the time spend together, listening to your stories, and engaging in debates on any topic. But I would like to extend my gratitude to all Agnes' family, for your kindness, embracing me as one of your own, and for treating me so well.

To my father Mariano, you have been a role model in terms of effort, hard work, dedication, but also honesty. Your guidance and occasional admonishments in my childhood were instrumental in shaping my character and helping me to strive to reach my goals. To my mother, Carme, you are the best mother anyone could have asked for. Your patience, tolerance, kindness, empathy and even too much solidarity, for which we have sometimes disagreed, have deeply influenced me. Not only this, but also your wit, irony and sarcasm, which I believe I've inherited, add a unique charm to our relation. You've taught me to think independently, debate, and express my thoughts, for which I am very thankful. To my brother David, your influence on my upbringing has also been immeasurable. Your early lessons in sharing (french fries) and your excellence in both sports and academics set a high bar that I aspired to reach during my all education. If you hadn't paved the way, maybe I wouldn't be here now. To you, Carol and the twins Ona and Nadia, I wish you all the best.

And finally, to my fiancée Agnès (Nessi), you encouraged me to embark in this journey, and you have been my driving forces behind my success during these five years. Your unwavering belief in me, even when I faltered, provided the motivation to persevere. You have been my rock, my confidant, and my source of inspiration. Your smile and jokes have had the remarkable ability to relieve my spirits, even in the most challenging situations. I have also learned a lot from you, from geology, fluid inclusions, mountains and their meteorology, but also to have an emotional relationship, listening, empathizing, learning, and trying to improve every day. As I move forward into the next chapter of our lives together, I want to express my deepest gratitude and dedicate this thesis to you as a symbol of our shared dreams and aspirations. Just as you have been with me during this academic journey, I know we will continue being together to support and uplift each other in our future endeavors.

To everyone I've crossed paths with: Each interaction, your knowledge, and your unique perspectives have contributed to who I am today and how far I've come. Thank you for being a part of my journey.

With my deepest gratitude,

Rubèn

Agraïments

Una tesi no es pot escriure sense la gent que tens al voltant. He conegut moltes persones fantàstiques al llarg del meu doctorat i el mínim que puc fer és reconèixer la seva ajuda.

Ernest, la teva paciència i ajuda durant el meu doctorat han estat inestimables. M'has guiat cap als temes de recerca més interessants i prometedors, i sempre m'has aportat molt bons comentaris i correccions a tots els articles, informes i finalment, a aquesta tesi. El teu suport al llarg d'aquests anys ha tingut un paper fonamental en la meva finalització amb èxit del doctorat, i per això, estic profundament agraït. Dimos, has estat com el meu segon supervisor no oficial. La nostra col·laboració, les discussions i les refutacions de les revisions dels articles han contribuït molt al meu creixement personal i professional. Sempre recordaré el teu consell de “Atacar els revisors en comptes de defensar-te d'ells”. A més a més, haig de reconèixer que va ser la teva passió per aquest camp i haver-te tingut com a professor a l'assignatura de Reconeixement de Patrons que van encendre el meu interès a dirigir la meva carrera professional cap a la Intel·ligència Artificial.

A tots els meus amics, però en especial a l'Ali i a l'Andrés. Hem estat junts des de l'inici del màster i ens hem fet costat en tot moment. Ali, la teva disposició per ajudar els altres és sorprenent i commovedora, des d'oferir-te per ajudar a tota la classe amb les assignatures de Seguretat i Informació la primera vegada que ens vam conèixer fins a passar hores revisant errors al codi amb mi. Al llarg de tot el doctorat, m'has ajudat molt, donant-me moltes idees, parlant de projectes esbojarrats (*Agronet*), jugant a squash per alliberar l'estrès o simplement passant un temps parlant de qualsevol cosa. Valoro tot el temps que hem passat junts i espero amb moltes ganes el dia que et pugui portar la muntanya, com porto insistent els darrers set anys. Andrés, els nostres debats sobre tota mena de temes, des de la consciència, la salut, l'esport o el veganisme, entre d'altres, juntament amb les nostres experiències compartides, han enriquit la meua vida de moltes maneres. Gràcies als dos pel vostre suport incondicional. Als tots els meus altres companys, amb els quals he passat molt de temps i aprecio molt, Sounak, Sanket, Sergi, Moha, Lei, Raul, Khanh, Pau, Adrià, Andrea, Alvin, Ayan, així com tots els estudiants que han vingut al CVC a fer una estada procedents de l'Índia, Itàlia i tot el món en general. També agraeixo als postdocs i professors dels quals he après coneixements molt: Lluís, Marçal, Josep, Alícia i Oriol. A més, he d'agrair també haver pogut gaudir d'un entorn de treball genial, on tothom que pot ajuda i que fomenta una fantàstica relació, incloent-hi també el personal d'administració i d'IT. Ha estat un privilegi ser part d'aquesta meravellosa comunitat.

Tot i tot el suport que he rebut al CVC, no hauria pogut completar el doctorat sense tota l'ajuda i l'escalfor de l'entorn a la meua vida personal. Angi i Xavier, la vostra càlida benvinguda em va fer sentir com a part de la família des del principi i sempre m'heu tractat com el vostre propi fill. També he après molt de vosaltres, vivint més a prop de la natura i adonant-me que els diners potser no són tan importants. Iaia Tere i tieta Marta, també m'heu acollit com si fos un altre net i nebot. Valoro molt tot el temps que hem passat junts, escoltant les vostres històries i debatent durant hores sobre qualsevol tema. Faig extensiu també el meu agraïment a tota la família de l'Agnès, per la vostra amabilitat i per haver-me tractat com un més de la família des del principi.

Al meu pare Mariano, has estat un model a seguir en termes d'esforç, treball dur i dedicació, però també d'honestetat i sinceritat. Els teus ensenyaments i els ocasionals càstigs quan em portava malament durant la meua infantesa han sigut fonamentals en la formació del meu caràcter i en ajudar-me a esforçar-me per assolir els meus objectius. A la meua mare Carme, ets la millor mare que ningú hagués pogut desitjar. La teva paciència, tolerància, amabilitat, empatia i fins i tot la teva solidaritat, encara que en excés en algunes ocasions, m'han influenciat profundament. No només això, sinó també el teu enginy, ironia i sarcasme, dels quals crec que n'he heretat una part, aporten un encant únic a la teva personalitat. M'has ensenyat a pensar de manera independent, debatre, saber expressar les meves idees alhora que respectar als altres, i per això estic molt agraït. Al meu germà, en David, la teva influència en la meua educació també ha estat incalculable. Les teves lliçons sobre compartir (patates fregides) i els teus bons resultats tant en els esports com en l'àmbit acadèmic van establir un llistó molt alt al qual sempre vaig aspirar a arribar durant tota la meua infantesa. Si tu no haguessis marcat el camí, potser jo ara no estaria aquí. A tu, a la Carol, i a les bessones Ona i Nadia, us desitjo tot el millor.

I finalment, a la meua promesa Agnès (Nessi), tu em vas encoratjar a embarcar-me en aquest viatge, i has estat la meua força impulsora durant aquests cinc anys. La teva confiança en mi, fins i tot quan a mi més em costava, m'ha donat la motivació per perseverar. Has estat la meua roca, la meua confident, i la meua font d'inspiració. El teu somriure i les teves bromes han tingut sempre una habilitat destacable per animar-me, fins i tot en les situacions més difícils. També he après molt de tu, des de geologia, incusions fluides, les muntanyes i la seva meteorologia, però també a tenir una relació, escoltant, empatitzant, aprenent i intentant millorar cada dia. Ara que ens encaminem cap al següent capítol de les nostres vides junts, vull expressar-te la meua més profunda gratitud. I així com has estat amb mi durant aquest viatge acadèmic, sé que continuarem junts estimant-nos, donant-nos suport i animant-nos l'un a l'altre en el futur que ens espera.

A tothom a qui he conegut pel camí: Cada interacció, el vostre coneixement i les vostres perspectives úniques han contribuït a ser qui sóc avui, i a arribar fins aquí. Gràcies per haver compartit una part d'aquest viatge amb mi.

Amb la meua més profunda gratitud,

Rubèn

Abstract

Visual Question Answering (VQA) is the task where given an image and a natural language question, the objective is to generate a natural language answer. At the intersection between computer vision and natural language processing, this task can be seen as a measure of image understanding capabilities, as it requires to reason about objects, actions, colors, positions, the relations between the different elements as well as commonsense reasoning, world knowledge, arithmetic skills and natural language understanding. However, even though the text present in the images conveys important semantically rich information that is explicit and not available in any other form, most VQA methods remained illiterate, largely ignoring the text despite its potential significance. In this thesis, we set out on a journey to bring reading capabilities to computer vision models applied to the VQA task, creating new datasets and methods that can read, reason and integrate the text with other visual cues in natural scene images and documents.

In **Chapter 3**, we address the combination of scene text with visual information to fully understand all the nuances of natural scene images. To achieve this objective, we define a new sub-task of VQA that requires reading the text in the image, and highlight the limitations of the current methods. In addition, we propose a new architecture that integrates both modalities and jointly reasons about textual and visual features.

In **Chapter 5**, we shift the domain of VQA with reading capabilities and apply it on scanned industry document images, providing a high-level end-purpose perspective to Document Understanding, which has been primarily focused on digitizing the document's contents and extracting key values without considering the ultimate purpose of the extracted information. For this, we create a dataset which requires methods to reason about the unique and challenging elements of documents, such as text, images, tables, graphs and complex layouts, to provide accurate answers in natural language. However, we observed that explicit visual features provide a slight contribution in the overall performance, since the main information is usually conveyed within the text and its position. In consequence, in **Chapter 6**, we propose VQA on infographic images, seeking for document images with more visually rich elements that require to fully exploit visual information in order to answer the questions. We show the performance gap of different methods when used over industry scanned and infographic images, and propose a new method that integrates the visual features in early stages, which allows the transformer architecture to exploit the visual features during the self-attention operation.

Instead, in **Chapter 7**, we apply VQA on a big collection of single-page documents, where the methods must find which documents are relevant to answer the question, and

provide the answer itself. Finally, in **Chapter 8**, mimicking real-world application problems where systems must process documents with multiple pages, we address the multipage document visual question answering task. We demonstrate the limitations of existing methods, including models specifically designed to process long sequences. To overcome these limitations, we propose a hierarchical architecture that can process long documents, answer questions, and provide the index of the page where the information to answer the question is located as an explainability measure.

Keywords – Vision and Language, Computer Vision, Pattern Recognition, Deep Learning, Scene Text Visual Question Answering, Document Visual Question Answering, Document Collections, Multipage Documents

Resum

Visual Question Answering (VQA) és la tasca on donada una imatge i una pregunta en llenguatge natural, l'objectiu és generar una resposta en llenguatge natural. En la intersecció entre la visió per computador i el processament del llenguatge natural, aquesta tasca es pot veure com una mesura de les capacitats de comprensió de les imatges, ja que requereix raonar sobre objectes, accions, colors, posicions, les relacions entre els diferents elements, així com emprar el sentit comú, coneixements globals, habilitats aritmètiques i comprensió del llenguatge natural. No obstant, tot i que el text present a les imatges conté informació semànticament important que és explícita i no disponible de cap altra forma, la majoria dels mètodes VQA són analfabets, ignorant en gran part el text de les imatges malgrat la seva potencial importància. En aquesta tesi, iniciem un viatge per aconseguir aportar capacitats de lectura als models de visió per computador aplicats a la tasca VQA, creant nous conjunts de dades i mètodes que puguin llegir, raonar i integrar el text amb altres indicis visuals en imatges d'escenes naturals i documents.

Al **Capítol 3**, adrecem el problema de com combinar el text en escenes naturals amb la informació visual per entendre completament tots els matisos de les imatges. Per a això, definim una nova subtasca de VQA que requereix llegir el text de la imatge, destacant les limitacions dels mètodes actuals i, en conseqüència, proposem una nova arquitectura que pot raonar sobre ambdues modalitats.

Al **Capítol 5**, canviem el domini de VQA amb capacitats de lectura i l'apliquem en imatges de documents escanejats, proporcionant una perspectiva final d'alt nivell al camp dedicat a intel·ligència sobre documents, els quals s'ha centrat principalment a digitalitzar el contingut d'aquests documents i extreure valors clau sense tenir en compte la finalitat per a la qual s'extreia la informació. Per a això, creem un conjunt de dades que requereix que els mètodes raonin sobre els elements únics que es poden trobar en documents, com ara text, taules, figures o gràfics disposats en complexos dissenys, per a aconseguir proporcionar respostes precises en llenguatge natural. No obstant això, vam observar que les característiques visuals explícites aporten una lleugera contribució al rendiment total, ja que la informació es troba principalment al mateix text i en la seva posició. En conseqüència, al **Capítol 6**, proposem aplicar VQA sobre imatges infogràfiques, buscant imatges de documents amb elements visualment més rics que requereixin explotar al màxim la informació visual per respondre les preguntes. En aquest capítol mostrem la diferència de rendiment de diferents mètodes quan s'utilitzen sobre imatges de documents escanejats o sobre imatges infogràfiques, i proposem un nou mètode que integra les característiques visuals en les primeres etapes, el qual permet que l'arquitectura del model exploti les

característiques visuals durant l'operació d'autoatenció.

En canvi, al **Capítol 7**, apliquem VQA a una gran col·lecció de documents d'una sola pàgina, on els mètodes han de trobar quins documents són rellevants per respondre la pregunta i proporcionar la resposta en si. Finalment, al **Capítol 8**, imitant els escenaris reals on els sistemes han de processar documents amb diverses pàgines, proposem utilitzar documents multipàgina per a VQA mostrant les limitacions dels mètodes ja existents. A més, proposem una arquitectura jeràrquica que pot processar documents llargs, respondre les preguntes plantejades i proporcionar l'índex de la pàgina on es troba la informació per respondre la pregunta com a mesura d'explicabilitat.

Paraules Clau – Visió i Llenguatge, Visió per Computador, Reconeixement de Patrons, Aprenentatge Profund, Resposta de Preguntes sobre Imatges amb Text, Resposta de Preguntes sobre Imatges de Documents, Col·leccions de Documents, Documents Multipàgina

Resumen

Visual Question Answering (VQA) es la tarea en la que dada una imagen y una pregunta en lenguaje natural, el objetivo es generar una respuesta en lenguaje natural. En la intersección entre la visión por computador y el procesamiento del lenguaje natural, esta tarea puede verse como una medida de las capacidades de comprensión de las imágenes, ya que requiere razonar sobre objetos, acciones, colores, posiciones, las relaciones entre los diferentes elementos, así como utilizar el sentido común, conocimientos globales, habilidades aritméticas y comprensión del lenguaje natural. Sin embargo, aunque el texto presente en las imágenes contiene información semánticamente importante que es explícita y no disponible en ninguna otra forma, la mayoría de los métodos de VQA son analfabetos, ignorando en gran parte el texto de las imágenes a pesar de su potencial importancia. En esta tesis, iniciamos un viaje para conseguir aportar capacidades de lectura a los modelos de visión por computador aplicados a la tarea VQA, creando nuevos conjuntos de datos y métodos que puedan leer, razonar e integrar el texto con otros indicios visuales en imágenes de escenas naturales y documentos.

En el **Capítulo 3**, dirigimos el problema de cómo combinar el texto en escenas naturales con la información visual para entender completamente todos los matices de las imágenes. Para ello, definimos una nueva subtask de VQA que requiere leer el texto de la imagen, destacando las limitaciones de los métodos actuales y, en consecuencia, proponemos una nueva arquitectura que puede razonar sobre ambas modalidades.

En el **Capítulo 5**, cambiamos el dominio de VQA con capacidades de lectura y lo aplicamos en imágenes de documentos escaneados, proporcionando una perspectiva final de alto nivel en el campo dedicado a la inteligencia sobre documentos, el cual se ha centrado principalmente en digitalizar el contenido de estos documentos y extraer valores clave sin tener en cuenta la finalidad para la que se extraía dicha información. Para ello, creamos un conjunto de datos que requiere que los métodos razonen sobre los elementos únicos que pueden encontrarse en documentos como texto, tablas, figuras o gráficos dispuestos en complejos diseños, para conseguir proporcionar respuestas precisas en lenguaje natural. Sin embargo, observamos que las características visuales explícitas aportan una ligera contribución al rendimiento total, puesto que la información se encuentra principalmente en el mismo texto y en su posición. Por ello, en el **Capítulo 6**, proponemos aplicar VQA sobre imágenes infográficas, buscando imágenes de documentos con elementos visualmente más ricos que requieran explotar al máximo la información visual para responder a las preguntas. En este capítulo mostramos la diferencia de rendimiento de diferentes métodos cuando se utilizan sobre imágenes de documentos escaneados o sobre imágenes

infográficas, y proponemos un nuevo método que integra las características visuales en las primeras etapas, que permite que la arquitectura del modelo explote las características visuales durante la operación de autoatención.

En cambio, en el **Capítulo 7**, aplicamos VQA a una gran colección de documentos de una sola página, donde los métodos deben encontrar qué documentos son relevantes para responder a la pregunta y proporcionar la respuesta en sí. Por último, en el **Capítulo 8**, imitando los escenarios reales donde los sistemas deben procesar documentos con varias páginas, proponemos utilizar documentos multipágina para VQA mostrando las limitaciones de los métodos ya existentes. Además, proponemos una arquitectura jerárquica que puede procesar documentos largos, responder a las preguntas planteadas y proporcionar el índice de la página donde se encuentra la información para responder a la pregunta como medida de explicabilidad.

Palabras Clave – Visión y Language, Visión por Computador, Reconocimiento de Patrones, Aprendizaje Profundo, Respuesta de Preguntas sobre Imágenes con Texto, Respuesta de Preguntas sobre Imágenes de Documentos, Colecciones de documentos, Documentos Multipágina

Contents

1	Introduction	1
1.1	Vision	2
1.2	Language	5
1.3	Scene Text	6
1.4	Document Understanding	7
1.5	Outline, Research Questions and Contributions	8
2	Background	11
2.1	Text Detection and Recognition	11
2.2	Word embeddings	13
2.3	Visual Question Answering	14
2.3.1	Datasets	14
2.3.2	Methods	17
I	Natural Scenes	19
3	Scene Text Visual Question Answering	21
3.1	Introduction	22
3.2	ST-VQA Dataset	24
3.2.1	Dataset creation	24
3.2.2	Analysis and Comparison with TextVQA	25
3.2.3	Tasks	28
3.2.4	Evaluation metric	30
3.3	Baselines	31
3.3.1	Baselines results	32
3.4	MMGrid Method	35
3.5	MMGrid results	38
3.6	Evolution of ST-VQA	41
3.7	Conclusions and Future Work	43

II Documents	45
4 Background	47
4.1 Document Understanding	47
4.2 Machine Reading Comprehension and Question Answering	48
5 Visual Question Answering on Single Page Documents	51
5.1 Introduction	52
5.2 SP-DocVQA Dataset	54
5.2.1 Dataset creation	54
5.2.2 Statistics and Analysis	56
5.3 Baselines	60
5.3.1 Heuristics and Upper Bounds	60
5.3.2 VQA Models	60
5.3.3 Reading Comprehension Models	61
5.4 Experiments	62
5.4.1 Experimental setup	62
5.4.2 Results	63
5.5 Evolution of VQA on Single Page Documents	71
5.6 Conclusion	72
6 Visual Question Answering on Infographic Images	75
6.1 Introduction	76
6.2 InfographicsVQA Dataset	78
6.2.1 Dataset creation	78
6.2.2 Question-answer types	78
6.2.3 Summary, statistics, and analysis	80
6.3 Baselines	82
6.3.1 Heuristic baselines and upper bounds	82
6.3.2 VQA Models	83
6.4 Experiments	85
6.4.1 Experimental setup	85
6.4.2 Results	86
6.5 Evolution of VQA on Infographic Images	96
6.6 Conclusion	97
III Document collections and multipage documents	99
7 Visual Question Answering on Document Collections	101
7.1 Introduction	102
7.2 DocCVQA Dataset	103
7.2.1 Dataset creation	103
7.2.2 Evaluation metrics	106
7.3 Baselines	107
7.3.1 Text spotting + QA	107

7.3.2	Database approach	108
7.4	Results	109
7.4.1	Evidences	109
7.4.2	Answers	110
7.5	Evolution of VQA on Document Collections	112
7.6	Conclusions and Future Work	113
8	Visual Question Answering on Multipage Documents	115
8.1	Introduction	116
8.2	MP-DocVQA Dataset	118
8.2.1	Dataset creation	118
8.2.2	Dataset statistics	120
8.3	Baselines	121
8.4	Hi-VT5	123
8.5	Experiments	126
8.5.1	Baseline results	127
8.5.2	Hi-VT5 results	128
8.6	Ablation study	132
8.7	Conclusions	134
9	Conclusions and Future Directions	135
9.1	Conclusions	135
9.2	New Directions	137
	List of Contributions	139
	Bibliography	141

List of Tables

2.1	State-of-the-art methods on Text Spotting in 2019 and 2023	12
3.1	Images and questions statistics of ST-VQA.	24
3.2	Baselines results on the ST-VQA dataset.	33
3.3	MMGrid performance.	39
3.4	Evolution of datasets in Scene Text VQA	41
5.1	Results of heuristic methods and upper bounds of SP-DocVQA	63
5.2	Performance of ST-VQA models on SP-DocVQA	64
5.3	Performance of BERT question answering model on SP-DocVQA.	65
6.1	Summary of VQA and Multimodal QA datasets.	77
6.2	InfographicsVQA dataset statistics.	80
6.3	Results of heuristic methods and upper bounds of InfographicsVQA	86
6.4	M4C results on InfographicsVQA.	87
6.5	BERT results on InfographicsVQA	87
6.6	LayoutLM results on InfographicsVQA	88
7.1	DocCVQA form fields analysis.	104
7.2	DocCVQA test set questions.	105
7.3	Baselines retrieval performance in DocCVQA	109
7.4	Baselines retrieval and answering results in DocCVQA	110
8.1	Comparison between MP-DocVQA and main DocVQA datasets.	118
8.2	MP-DocVQA dataset statistics along its construction process	119
8.3	MP-DocVQA methods hyperparameters.	126
8.4	Baselines and Hi-VT5 results on MP-DocVQA dataset	127
8.5	Hi-VT5 ablation study	132
8.6	Results of Hi-VT5 with different [PAGE] tokens.	133
8.7	Hi-VT5 results when trained with different number of document pages. . .	133

List of Figures

1.1	CNN learned features visualization.	3
1.2	Example of two Computer Vision Neural Networks	4
1.3	Relevance of Scene Text for different tasks.	6
2.1	VQA Datasets examples.	16
3.1	Scene-Text VQA examples.	23
3.2	Percentage of questions and answers per number of words.	26
3.3	Starting 4-grams of ST-VQA questions.	26
3.4	Answer variability per questions in ST-VQA.	27
3.5	ST-VQA proposed tasks examples.	29
3.6	ANLS and Accuracy metrics comparison.	30
3.7	Results per question type of baseline methods in ST-VQA.	34
3.8	Baselines qualitative results on ST-VQA.	35
3.9	MMGrid model architecture.	36
3.10	MMGrid cell assignment of the OCR words	37
3.11	Qualitative results and attention masks visualization of MMGrid in ST-VQA	40
3.12	Evolution of method's performance in ST-VQA dataset.	42
5.1	SP-DocVQA example.	53
5.2	SP-DocVQA source documents statistics.	54
5.3	SP-DocVQA document's industry source.	55
5.4	SP-DocVQA question types.	56
5.5	SP-DocVQA question statistics.	57
5.6	SP-DocVQA answer statistics.	58
5.7	SP-DocVQA recognized text statistics.	59
5.8	Baselines performance per answer type.	65
5.9	SP-DocVQA qualitative results 1	66
5.10	SP-DocVQA qualitative results 2	67
5.11	SP-DocVQA qualitative results 3	68
5.12	SP-DocVQA qualitative results 4	69
5.13	SP-DocVQA qualitative results 5	70
5.14	Evolution of method's performance in SP-DocVQA dataset.	71

6.1	InfographicsVQA example.	76
6.2	InfographicsVQA percentage of questions per types.	79
6.3	Staring 4-grams of questions in InfographicsVQA.	80
6.4	InfographicsVQA word clouds.	81
6.5	InfographicsVQA images layout statistics.	81
6.6	LayoutLM based model architecture for InfographicsVQA	83
6.7	Performance of baselines for different QA types.	89
6.8	InfographicsVQA qualitative results 1	89
6.9	InfographicsVQA qualitative results 2	90
6.10	InfographicsVQA qualitative results 3	91
6.11	InfographicsVQA qualitative results 4	92
6.12	InfographicsVQA qualitative results 5	93
6.13	InfographicsVQA qualitative results 6	94
6.14	InfographicsVQA qualitative results 7	95
6.15	Evolution of method’s performance in InfographicsVQA dataset.	96
7.1	DocCVQA task examples.	103
7.2	Number of relevant documents per question of DocCVQA.	105
7.3	Baselines evidence retrieval per question on the DocCVQA test set.	109
7.4	Baselines answering performance per question on the DocCVQA test set.	111
7.5	Evolution of method’s performance in DocCVQA dataset	112
8.1	MP-DocVQA task example.	116
8.2	MP-DocVQA dataset statistics.	120
8.3	Architecture of Hi-VT5 method	123
8.4	Methods’ performance by answer page position.	128
8.5	Answer-page retrieval correlation of Hi-VT5	129
8.6	Hi-VT5 qualitative examples in MP-DocVQA.	130
8.7	Visualization of the Hi-VT5 attention scores	131

Chapter 1

Introduction

Humans constantly perceive and communicate with their environment through the use of their five senses: sight, hearing, touch, taste, and smell. Perception is the process of receiving, organizing, and interpreting sensory information from the environment. In humans, this is conducted by our abilities to see, hear, touch, taste, and smell, as well as higher-level cognitive processes such as attention and memory, which allow us to recognize and identify objects, events, and patterns in the world around us. Understanding goes beyond mere recognition and involves the higher-level interpretation and integration of perceived information into a coherent and meaningful whole. It involves not only recognizing and identifying individual objects or events, but also grasping their relationships, significance, and context. Understanding involves the use of knowledge, concepts, and prior experiences to make sense of new information and to draw conclusions and make judgments about it. In addition, we also have developed complex systems of communication, including language and nonverbal cues, which allow us to exchange information. By combining the understanding of the environment and the communication of the processed information, humans have been able to develop a rich and nuanced understanding of the world. Because it is through communication that individuals share information, ideas, and perspectives, allowing to construct a shared understanding of the world, which in turn, can influence the way others perceive the world around them.

While perception in computer science is carried out by specific devices such as cameras, microphones or sensors, reasoning about them is a complex task studied in Artificial Intelligence (AI) research, lately performed through the use of Machine Learning (ML) algorithms. These methods can learn to recognize patterns and relationships in the data, and make predictions or decisions based on this learning. More specifically, Computer Vision (CV) is the AI field dedicated to interpret and understand visual data from the world around us, analyzing information of images from nature, cities, buildings, roads, or any other environment that exists in the real world. In addition, the Document Image Analysis and Recognition (DIAR) is a specific domain within computer vision that focuses on processing and analyzing images of man-made artifacts such as books, documents, posters, radiographies or any other written or printed material, specially used to convey information or ideas where text is specially relevant.

On the other hand, the study of the communication is carried out by the Natural Language Processing (NLP) field, which models the interaction between computers and human (natural) languages, focusing on the development of algorithms and models that can process, analyze, interpret and generate natural language in order to facilitate human-computer communication.

In the intersection between vision and language, Visual Question Answering (VQA) is a task where an AI agent takes as input an image and a natural language question about the image in a free-form, open-ended format, and must provide an accurate natural-language answer as the output. These questions may include identifying specific objects or attributes in an image, recognizing activities or emotions, determining the presence of certain characteristics, or using background knowledge to understand the scene. Therefore, VQA can be seen as a means of assessing the agent's comprehension ability about the images, and therefore, about the world and situations they depict.

However, there is crucial information that has been long ignored in VQA models. The text present in the images conveys unique information that is explicit and not available in any other form within the scene. Particularly in man-made environments, it is essential in order to perform most everyday tasks like making a purchase, using public transportation, finding a place in the city, getting an appointment, or checking whether a store is open or not among others. Moreover, this significance is further accentuated when dealing with documents like reading a restaurant's menu, reviewing an invoice or an application form, or understanding product manuals.

In this thesis, we set out on a journey to bring reading capabilities to CV models applied to the VQA task, creating new datasets and methods that can read and integrate the text with other visual cues in order to reason about all the contents of images in natural scenes and documents, and capture all the nuances of the images. In the following subsections, we present an overview of the diverse modalities used in the VQA task, essential for contextualizing this thesis.

1.1 Vision

According to Marr [141], vision is the process of discovering from images what is present in the world, and where it is. In humans, the light coming from the environment is captured by the photoreceptor cells in the retina of the eye, where it is converted into electrical signals that are sent to the brain for further processing and interpretation, constructing depictions of color, shape, edge, and locations of objects and their interactions with the environment. Researchers, have taken inspiration from the human biological process, dividing the perception and interpretation of the visual information into two streams. The perception is conducted by cameras that are sensitive to the electromagnetic radiation of the light through photodiodes, which captures the light, and converts it into electric signals divided into three main components, red, green and blue (RGB), which are later represented as color channels. The combination of such channels will create different colors and tonalities, and the amount of sensors used during the light capture will determine

the number of pixels of the final image. Hence, images in computers encode different light intensities of a pixel as a 2-dimensional matrix for each color channel. Dealing with this way of representing an image provides several benefits to visualizing, compressing and retrieving large-scale imaging data. However, such representation is hardly useful to obtain semantically rich information (present objects, interactions, spatial relations, text, etc.) by standard programming approaches.

Therefore, computer vision approaches have benefited from machine learning methods to develop algorithms that are able to see and interpret digital imagery. In the early days of computer vision, researchers used handcrafted methods [13, 132] by manually designing heuristics and algorithms to extract information from images. However, after the success of Neural Networks (NN) and Convolutional Neural Networks (CNN) [106, 110], they have been the standard approach in order to automatically learn representations from the images. CNNs consist of a set of convolutional layers with multiple learnable filters, called kernels, that are used to scan the input image. Each filter is responsible for detecting a specific feature in the image, and the result of the convolution operation is a feature map, which encodes the presence of the specific feature in different parts of the image. While early convolutional layers of a CNN detect low-level features such as edges and textures (Figure 1.1 left), the deeper layers, which combine the information from multiple lower-level features, detect more complex patterns such as entire objects or parts of them (Figure 1.1 right). Later advancements on top of CNNs that overcame problems such as vanishing gradients and improved the computational efficiency, allowed deeper and bigger networks [64, 65, 171]. More recently, Transformer based [211] approaches [24, 44] have been employed with success to achieve better results or obtain compact image embeddings. In a vision transformer architecture, the image is first divided into a fixed-size grid of non-overlapping patches, and then each patch is transformed into a feature vector using a linear embedding or a CNN. These feature vectors are then processed by the transformer layers, which perform self-attention and feed-forward operations to capture the relationships between the patches in the image. At the end, the output of the transformer layers is processed by a linear layer to produce the final prediction. Alternatively, the image representation of the last transformer layer can be used to represent the image in multimodal or retrieval tasks. However, these approaches require a high amount of annotated data and computational resources.




Figure 1.1: Visualization of features learned by different kernels from earlier (left) to later (right) convolutional layers.

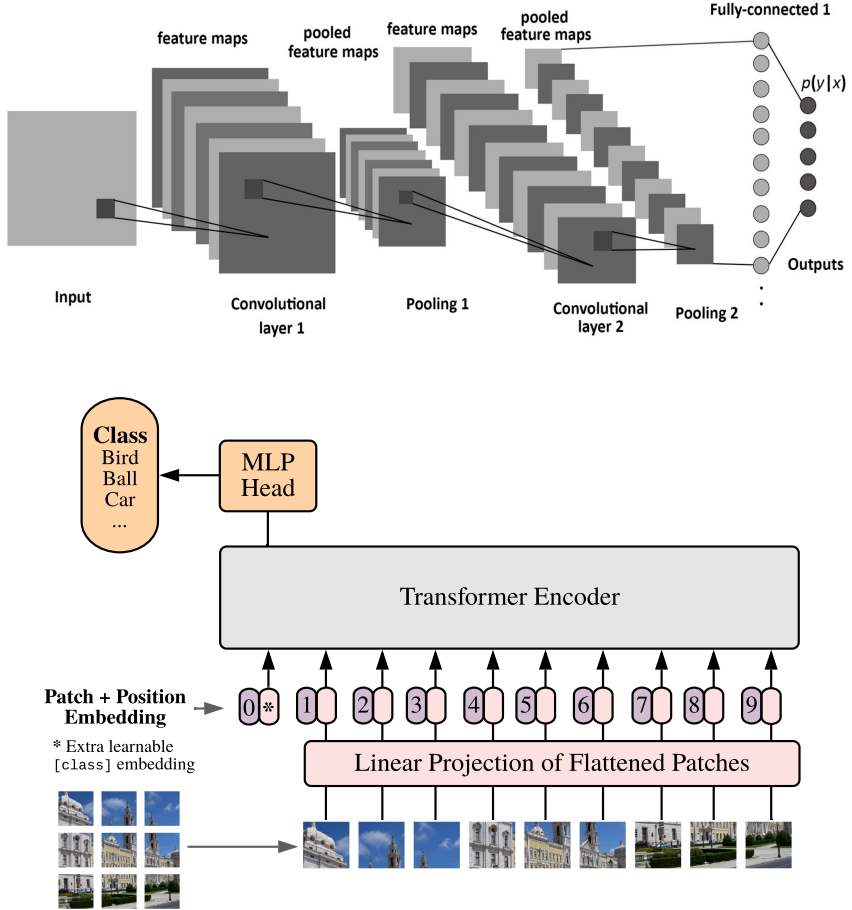


Figure 1.2: Example of two Computer Vision Neural Networks. At the top, a Convolutional Neural Network (CNN) [110] where convolutional and pooling layers alternate to extract a series of features through kernels, and reduce their dimension, while maintaining the most important information. In the end, a set of linear layers project these features to the final dimension that will depend on the task, and will enable making a prediction. At the bottom, a Vision Transformer (ViT) [44] network, where each patch in the image is embedded by a linear projection and a positional embedding is used to determine the order of the patches in the image. Then, in the Transformer encoder all patches attend to all the other patches through the multiple self-attention layers. At the end, a Multi Layer Perceptron (MLP) is used to project the contextualized features into the final dimension, and perform the prediction.

1.2 Language

Chomsky [34] defines language as the ability of speakers to comprehend and construct grammatical sentences. In humans, language is processed through a complex interplay of cognitive processes, including perception, comprehension, and generation. In computers, it is captured through textual or spoken input through a keyboard or microphone, where it's converted into a format that computers can process. This input may include words, sentences, or even entire documents. Natural Language Processing (NLP) is the field that studies and designs algorithms capable of understanding and interpreting human languages. The initial step for processing it involves transforming raw text into a structured representation that can be analyzed and manipulated with techniques such as tokenization, part-of-speech tagging, and syntactic parsing to break down the language into its fundamental units and uncover its grammatical structure. However, dealing with language in computational systems presents its own set of challenges. Unlike images that are represented as matrices, text is inherently sequential, possesses a hierarchical structure and exhibits nuances, such as ambiguity, context-dependency, and cultural variations, making it a complex and dynamic domain, where standard programming approaches often fall short in capturing the richness and subtleties of human language.

Even though it is possible to represent words with one-hot vectors by encoding each word in a dictionary, such an approach would result in a very sparse and fixed dictionary unable to incorporate new words. In consequence, earlier methods [19, 87, 109] attempted to model the language by building handcrafted features in a hierarchical manner, modelling the co-occurring statistics of the words. Later approaches [147, 159] relied on neural networks to encode each word with a compact, low-dimensional vector embedding, which allowed clustering the words according to different features such as the words' morphology, their syntactic role or their semantics. However, these approaches required large vocabulary dictionaries to encode all possible words. To alleviate this problem, later approaches [20, 42] decomposed each word into its most common n -grams, which allowed to represent more diverse vocabulary with smaller dictionaries. Such compact embeddings were also able to capture the semantics or morphology of the words, clustering similar words together and providing special mathematic properties.

In order to encode longer sequences, such as captions or long paragraphs, the standard approach consisted on directly employing single-word pre-trained word embeddings to be later fed into a Recurrent Neural Network (RNN) [33, 69, 174]. However, NLP has recently had a major breakthrough with the introduction of Transformers [211], which consist of a set of multi-headed attention layers, each on top of the other, where the input is decomposed into queries, keys and values, that are used to compute the attention in an all vs all fashion. Such architecture has yielded to significant improvement gains compared to RNNs, while at the same time avoids problems such as vanishing gradients and allows longer sequences to be modeled.

1.3 Scene Text

Scene text is a type of unstructured text that is present in natural images or video frames. It provides valuable and explicit semantic information that is essential for a complete interpretation of the world representation. In particular, scene text is abundant in urban environments, where it can be found in signs, billboards, license plates, street names, advertisements or storefronts among others. Therefore, without considering the scene text, it is not possible to provide an answer or a detailed description of the content within an image. Instead, by accurately integrating scene text with visual cues, a richer and more nuanced representation can be achieved, leading to more accurate results in a wide variety of tasks such as image classification, image captioning, or visual question answering as demonstrated in Figure 1.3.



Image classification and retrieval
Building

Visual Question Answering

Q: What is the name of the restaurant?

A: UNKNOWN

Image Captioning

A man walking down the street



Image classification and retrieval
Hot Dog Restaurant

Visual Question Answering

Q: What is the name of the restaurant?

A: Essie's Original

Image Captioning

A man walking in front of Essie's hot-dog shop.

Figure 1.3: Scene text provides explicit information essential to fully understand the scenes, making it crucial for achieving accurate results in tasks such as image classification, visual question answering and image captioning.

Nevertheless, detecting and recognizing scene text is a challenging problem due to the wide variety of fonts, colors, orientations, and sizes that text can appear in, the background complexity and inference factors such as noise, blur, distortions, and occlusions, all of which contribute to the difficulty of the task. The task can be tackled by performing text

detection and recognition in two sub-tasks, first detecting and localizing the text instances in the image, and latter using a recognition system to transcribe and convert the content of the detected text regions into linguistic symbols. The other approach is to tackle the task in an end-to-end fashion, where the system performs both text detection and recognition in one unified pipeline.

In this thesis, rather than focusing on the text detection and recognition task itself, we will make use of already existing methods to leverage the semantics of scene text in order to capture all the nuances of the image.

1.4 Document Understanding

Documents are a rich source of information, often containing a significant amount of data within a relatively small amount of space. They play a critical role in our daily live, and specially in various fields such as government, education, research or many business like banks and insurance companies. Therefore, the ability to extract and understand the information contained within documents is essential to make informed decisions and take effective actions. However, automatic information extraction can be challenging, since it can be present in a wide range of forms including text, images, tables or graphs. This is particularly true in specialized fields such as legal or technical studies, where the information may be complex in nature. Furthermore, scanned images may be subject to degradation due to factors such as image binarization, poor scan quality, or deterioration of the original documents. In addition, the intricate layout of documents poses additional challenges for computers in effectively interpreting all the information contained within, and even determining the correct reading order may prove difficult.

Most works have focused on structured or semistructured documents such as forms, invoices, receipts, passports, ID cards, e-mails or contracts among others. In the earlier years, the research was mostly focused on text recognition, initially dealing with printed text (OCR systems) and then with handwritten text. Later works [41, 176] proposed systems that integrate different tasks such as document classification and information extraction to address specific business requirements. These systems relied on predefined rules and required the creation of specific templates for each new document type. Later on, learning-based methods [36, 157] allowed to automatically classify the type of document and identify relevant fields of information without predefined templates. Recent advances in deep learning [126, 228, 239] leverage natural language processing, visual feature extraction and graph-based representations in order to achieve a global view of the document that take into account word semantics and visual layout in the process of information extraction. The integration of all these modalities, provide an excellent opportunity to extend the capabilities of VQA models to documents, opening new avenues for research and innovation in this field.

1.5 Outline, Research Questions and Contributions

The thesis is divided in 3 parts, all focusing on applying Visual Question Answering that requires to integrate recognized text from the images in different domains and settings.

The first part is dedicated to VQA on natural scenes, where the scene text is essential to answer the questions. For this, first we provide the context and previous work on generic VQA methods (chapter 2), and then, propose a dataset and method tackling the task (chapter 3). The second part is focused on applying VQA on document images, we first provide the context of document understanding tasks and question answering methods (chapter 4), and explore VQA on single page document images on two different types of documents, industry scanned documents (chapter 5) and infographic images (chapter 6). Finally, the last part focuses on performing VQA on documents considering multiple images, a collection of multiple single page documents (chapter 7), and multipage documents (8). At the end of the thesis, we also present in chapter 9 the conclusions of the thesis and the future directions.

All the datasets and code created during the thesis have been made publicly available in the Robust Reading Competition Platform¹, where all the challenges [17, 205] held accompanying the different tasks have been hosted. In this platform, researchers can evaluate their methods against a public test dataset where we have a public leaderboard with tables of quantitative and qualitative results of submitted methods, which helps to easily compare different behaviors.

In the following paragraphs, we provide the reader with a summary of the chapters included in this thesis, with the research questions that motivated each work.

Part I – VQA on Natural Scenes

Chapter 3 - Scene Text Visual Question Answering

Research Question 1: Can VQA models integrate reading capabilities and leverage the text in the image to answer the questions?

Research Question 2: What type of data do we need to make VQA models literate?

We observe that VQA models are illiterate, incapable of answering any questions that require reading the text. In response, we create a new dataset called Scene-Text Visual Question Answering (ST-VQA), which emphasizes the importance of properly leveraging the high-level semantic information present in the form of scene text within images. We also provide several baselines, demonstrate the impact of using textual features and propose a new method that can point the answer directly to the scene text in the image by attending to multimodal grid features, allowing it to reason jointly about textual and visual modalities.

¹<https://rrc.cvc.uab.es/>

Part II – VQA on Documents

Chapter 5 - Visual Question Answering on Single Page Documents

Research Question 1: Can VQA task be used to guide information extraction on documents, providing the queries and the answers in natural language form?

Research Question 2: What are the unique challenges in VQA on documents that are not present in natural scenes? Can ST-VQA methods be applied on documents where textual information is largely predominant and relevant in documents?

Document understanding has focused on specific methods and pipelines for each information extraction task. Instead, we propose to use VQA on documents as a high-level task focused on the ultimate purpose the information will be used for, using VQA as a generic document understanding and information extraction interface by asking questions over the scanned document images. In this case, the information extraction is conditioned on a natural language question, and a natural language answer is expected as the output. To achieve this objective, we create the dataset Single Page Document Visual Question Answering (SP-DocVQA) where the questions are posed over industry scanned documents. Then, we use ST-VQA and NLP QA methods on the new dataset and demonstrate their limitations to handle complex document layouts, tables, and figures, as well as large amounts of text, which is often much greater than in natural scenes. This highlights the need for specific methods that can effectively tackle the unique challenges presented by documents.

Chapter 6 - Visual Question Answering on Infographic Images

Research Question 1: How can we leverage visual information in document images?

Research Question 2: How can VQA models perform complex reasoning integrating visual, textual and spatial modalities?

Due to the low relevance of explicit visual information on industry scanned documents for VQA, we consider to look for more complex images to explore the limits of current methods. For this, we use infographic images because they contain more complex and visual elements such as icons and graphics, and the color, shapes and visual appearance of such elements are, more often than not, visual cues that help linking different elements referring to the same concept. Then, we create a challenging dataset named InfographicsVQA that requires the methods to model spatial, visual and semantic information together. We provide several baselines to show that even the methods that use spatial and visual modalities perform poorly when compared against their performance on industry scanned documents, which highlights the need for a specific architecture that jointly reasons about all modalities. Hence, we propose a transformer method that integrates visual features in early stages to leverage the visual information during the self-attention operation.

Part III – VQA on Documents with multiple images

Chapter 7 - Visual Question Answering on Document Collections

Research Question 1: Is it possible to extract information from a collection of documents using a Visual Question Answering task as a natural language interface?

Research Question 2: Can Document Visual Question Answering methods process a collection of single-page documents?

While research on VQA on single page scanned documents progress rapidly, we consider if it is possible to process multiple images for each question-answer pair. Inspired by real applications where multiple documents are analyzed at the same time to extract aggregated information, we propose to recreate this scenario in Visual Question Answering. Due to the big size of the document collection, this task is formulated as a retrieval-answering task, which involves a two-stage process. First, the relevant documents are retrieved based on the given question, and then the question is answered using the retrieved documents.

Chapter 8 - Visual Question Answering on Multipage Documents

Research Question 1: Can Document Visual Question Answering methods process multiple images simultaneously?

Research Question 2: Is it possible to ascertain that a VQA model not only produces accurate responses, but also supplies an explanation for the reasoning behind its answers?

In this chapter, we apply VQA on multipage documents, in contrast to the previous collections of single page documents. This requires smaller sets of images and with this task we seek to model the interplay between the different pages in a single stage. Thus, we demonstrate the limitations of different DocVQA and QA methods for long sequences of tokens, and create a hierarchical method that can process long documents in a single step, processing all the pages simultaneously and providing the answer to the question. In addition, the model also includes a module to identify the location where the information to answer the question is found, enabling to identify the answers that have been inferred from the input data, and those derived from prior bias.

Chapter 2

Background

2.1 Text Detection and Recognition

The task of text detection and recognition in natural images sets the starting point for generic VQA systems to start integrating textual cues and move towards VQA systems with complete scene understanding. The most common approach in the reading systems community consists of two steps, text detection and recognition.

In text detection, several methods [121–123, 136, 180, 246] adapt object detection methods by modifying the region proposal and bounding box regression modules to localize text instances. An alternative approach is to predict sub-text components (pixels [39, 220], characters [8] or medium-granularity components [180, 203, 242]) to finally assemble them into a text instance. While sub-text component-based methods may face increased complexity in assembling the components into complete instances, which can be computationally expensive and more challenging than object detection-based methods for direct text localization, they offer the advantage of providing fine-grained information at the character-level or sub-component-level attributes. In addition, the flexibility of these methods allows them to better adapt to curved, long and oriented text.

In text recognition, the input consists of a cropped text instance which contains only one word. Two major techniques for text recognition are the Connectionist Temporal Classification (CTC) [58], and the encoder-decoder paradigm [195]. Both techniques use CNNs to encode the images into feature spaces but differ on how they decode the text. In CTC based methods [66, 181, 182] the input images are viewed as a sequence of vertical pixel frames, and the network outputs a prediction per each frame indicating the probability of all possible per-frame prediction for a text string by stacking RNNs on top of CNNs. In contrast, the encoder-decoder methods for text recognition [9, 31, 111, 130] uses an RNN to read an input sequence and pass its final latent state to a decoder RNN, generating the output in an autoregressive way. The main advantage of the encoder-decoder framework is that it provides outputs of variable lengths, which is particularly

useful when recognizing text of different lengths.

Recently, many methods [21, 23, 67, 127, 135, 163, 246] have opted to perform both text detection and recognition tasks end-to-end, which is known as text spotting. These methods are considered more efficient and robust, and often achieve better results by leveraging the context and layout of the entire image when recognizing text. However, these are considered as context-free methods because they lack of the understanding of the text semantics, and consequently they are less robust against occluded or incomplete text. In contrast, the current trend [12, 37, 47, 118, 152] has become to use language models to interpret the contextual scene text to correct possible errors produced by poor visual conditions on particular regions of the images. Moreover, along with language models, the recent use of Transformers for text spotting has allowed the methods [241, 243] to jointly perform curved text instance detection and recognition without the need of any heuristics-driven post-processing procedures, nor Region-of-Interest operations. Furthermore, Vision-Language models pre-trained on webscale data like CLIP [164] or ALIGN [84] has shown remarkable zero-shot capacity with strong scene text perception abilities, and can robustly perceive irregular text in natural images. In consequence, this makes them a good choice as the backbone for some methods [243], providing a superior starting baseline, and allowing them to achieve the best current results. There are several comprehensive surveys [30, 131] that cover the topic of Scene Text detection and recognition topic and provide further details on every task, dataset and method.

Dataset	Year	Best method 2019		Best method 2023	
		Name	Acc.	Name	Acc.
SVT [216]	2011	ASTER [183]	93.6	CLIP4STR [243]	98.3
IIIT-5K [148]	2012	ASTER [183]	93.4	CLIP4STR [243]	99.2
ICDAR-13 [94]	2013	FAN [31]	93.3	ParSeq [12]	98.4
SVTP [160]	2013	ASTER [183]	78.5	CLIP4STR [243]	97.2
CUTE [172]	2014	ASTER [183]	79.5	CLIP4STR [243]	99.3
ICDAR-15 [93]	2015	ASTER [183]	76.1	CLIP4STR [243]	91.4

Table 2.1: State-of-the-art Text Spotting in 2019 (at the start of the Ph.D.) and Current Performance in 2023. The performance of Text Spotting methods in 2019, especially on easy rectangular non-occluded text, allowed to glimpse the potential of utilizing text in images for higher-level tasks.

As shown in Table 2.1, state-of-the-art methods at the outset of this thesis were achieving an accurate detection and recognition performance of horizontal text on well-known datasets like SVT [216], IIIT-5K [148], and ICDAR-13 [94], with an accuracy rate exceeding 90%. While further improvements were still needed for text in arbitrary orientations and shapes as indicated by the performance on datasets like CUTE [172], ICDAR-15 [93], and SVTP [160], the maturity of the field and the promising results motivated the use of recognized text in complex tasks like Visual Question Answering.

2.2 Word embeddings

In order to process the text, either sourced from a machine-readable format or recognized text extracted from an image, it must be encoded as a vector. Even though it is possible to represent the words with a one-hot vector encoding the position of each word in a fixed dictionary, such approach would lead to a very sparse and fixed representation incapable of incorporating new words. Moreover, this representation would not account for the meaning, semantics, or morphology of the words, and consequently, a model using one-hot vectors would need to learn the meaning and relationships of the words from scratch, which can be challenging and inefficient. Instead, word embeddings enable word vectors to capture similarities and relationships between words that can be beneficial for text recognition tasks, such as correcting errors and handling misspellings, or in Visual Question Answering where the semantic and relationship between the question and the recognized words are crucial to answer the questions. Nevertheless, there are different word embeddings that aim to capture different properties.

PHOC [3] proposed a more refined use of binary vectors, encoding if a specific character appears in a particular spatial region of the string, creating a Pyramidal Histogram Of Characters at multiple levels. This compact representation doesn't require to be trained, captures the morphology of the words, and it does not require a fixed vocabulary. On the other hand, it doesn't encode any semantic information. Other approaches rely on Neural Networks and usually represent words as compact low-dimensional vector embeddings. In Word2Vec [147], the main idea is to teach a network to predict the most probable words that fit a common context given by the surrounding words (skip-gram) or the most likely neighbors of a word (contiguous bag-of-words). Therefore, if two words share the same context, they should end up having similar embeddings. Instead, GloVe [159] encodes the word semantics based on the ratios of word-to-word co-occurrence probabilities, and is therefore trained on a global word-to-word matrix of global co-occurrence by minimizing the Mean Square Error between the dot product of the word embeddings and their probability of co-occurrence.

However, such embeddings are designed to work with single tokens, each representing a different word. Consequently, these embeddings are equally encoded within a fixed dictionary that, although its vectors convey semantic meaning, is still fixed and can't deal with out-of-vocabulary words. To alleviate this problem, FastText [20] proposes to create the vocabulary with the most common n-grams [178], encoding each word with overlapping sub-word embeddings. It is trained with the skip-gram strategy to exploit the semantics and morphology of the words, while dealing with out-of-vocabulary words. The resulting pretrained embeddings acquire interesting mathematical properties such as:

$$\begin{aligned} king - man + woman &\approx queen \\ Paris - France + Germany &\approx Berlin \end{aligned}$$

thus, allowing the operation of algorithms in semantic spaces that allow models to better work with high-level concepts.

BERT [42] leverages Transformers’ self-attention to learn contextualized word embeddings. Unlike previous word embeddings that assign fixed vector representations to words regardless of their context, BERT takes into account the surrounding words and their order within a sentence to generate context-dependent word embeddings. Moreover, it introduced the concept of self-supervised pretraining tasks to learn meaningful representations from large amounts of unlabeled text data, to be latter fine-tuned on the downstream task with slight or no changes in the architecture. More specifically, it introduced the pretraining tasks of Masked Language Modeling and Next Sentence Prediction. In the former it predicts a missing word within a sentence, enhancing the understanding of word dependencies, and in the latter, it facilitates its application in tasks requiring understanding of sentence-level relationships. Moreover, besides the word-level representations, BERT introduces a compact representation of the whole sentence using a special token called [CLS], which is inserted at the beginning of the input sentence and allows to represent the whole sentence without the need of any Recurrent Neural Network. The success of this Transformer-based embedding led to a vast amount of new variations [51, 108, 115, 128, 165, 166, 175], based on different pre-training tasks and the use of more and varied datasets.

At the intersection between Vision and Language, Visual Question Answering requires to effectively model the relations between the words and thus, using the appropriate word embeddings is of crucial importance.

2.3 Visual Question Answering

Visual Question Answering (VQA) aims to come up with an answer to a given natural language question about the image. Since its introduction, VQA has received a lot of attention from the Computer Vision community [2, 6, 50, 57, 76, 86, 137, 170, 235]. In the following subsections, we first provide an overview of the VQA datasets, including key aspects such as their size, question-answer types, image diversity or specific challenges conveyed. Secondly, we describe previous VQA methods that don’t consider the recognized text in the scenes, focusing on the most popular approaches.

2.3.1 Datasets


The DAQUAR [137] dataset was the first toy-sized QA benchmark built upon indoor scene RGB-D images, but the limited variety of questions and size, reduced its effectiveness to assess real the method’s VQA capabilities. VQA (or VQA 1.0) [6] is considered the first modern VQA dataset due to its size and variability. The dataset contains both open-ended and multiple-choice questions for real scenes and clip-art images, which compounds two different subsets of the dataset. The dataset presented some biases. Several questions could be answered correctly without considering the images (binary questions represent 38% of the dataset, and 59% of them are answered with ‘Yes’). Moreover, because the images are sourced from MS COCO [124], many sport related questions can

be answered with ‘tennis’ or ‘baseball’, and questions asking for animals are usually answered with ‘cat’ and less frequently ‘dog’. VQA 2.0 [57] is an evolution of VQA 1.0 which balanced the dataset by collecting complementary images such that every question in the new dataset is associated with more than one image, resulting in two different answers to the question.

CLEVR [86] is a collection of 100,000 synthetic images of 3D shapes such as spheres and cylinders. It is built to contain minimal biases and has detailed annotations describing the kind of reasoning each question requires. This enables performing an in-depth analysis of the methods by slicing the dataset along different axes (question type, relationship type, question topology, etc.), which allows a better understanding of the reasoning capabilities of VQA system. Visual Genome [104] is a large-scale multitask image understanding dataset and knowledge base that combines images, structured metadata and natural language descriptions. The dataset contains over 108,077 images, each annotated with region-level objects and attributes, spatial relationships between objects, and free-form textual descriptions. It is the biggest question answering dataset with 1,773,258 question answer pairs (excluding automatically generated QA datasets). Moreover, the questions start with one of the six Ws (what, where, when, who, why, and how), which makes the dataset offer a considerable coverage of question types, ranging from basic perceptual tasks (e.g. recognizing objects and scenes) to complex common sense reasoning (e.g. inferring motivations of people and causality of events).


VizWiz [60] differs from the rest of the VQA datasets because the images were captured through a real application by blind photographers, and they often have poor quality with problems such as blur, inadequate lighting, finger covering the lens or poor framing of the image content. This actually turned $\sim 28\%$ of the questions to be deemed unanswerable. Moreover, the questions were originally spoken, and they look more conversational. Interestingly, regarding the topic of this thesis, a recent study [237] on VizWiz shows that text recognition is the second most common skill required to answer the visual questions. Specifically, the most common skill is object recognition (required in 93.25% of the questions), followed by text recognition (45.33%), color recognition (22.06%), and counting (1.49%).

Other works have focused on introducing new reasoning requirements, such as multilingual questions in FM-IQA [50], the need of prior knowledge to answer questions that can’t be answered directly from the image in KB-VQA [218], FVQA [219], and OK-VQA [140], arithmetic skills to solve complex counting questions in HowManyQA [210], and TallyQA [1], or the use of multiple sources of information in MovieQA [201], where the inputs to answer a question are the movie video clips, the plot synopses, the script, the subtitles, and the descriptive video service for visually impaired people. One work specially interesting regarding the third part of the thesis (VQA on Multipage Documents) is ISVQA [10], where the questions are posed on indoor and outdoor scenes depicted by a set of natural images. Within this set of images, some of them depict overlapping parts of the scene, and consequently the methods must learn to differentiate between a new perspective of an already seen region from a region not seen before.



Q: There is a rubber cube in front of the big cylinder in front of the big brown matte thing; what is its size?

A: Small



Q: What number of purple shiny things are there?

A: 1




Q: What is the mustache made of?

A: Bananas



Q: Is this a vegetarian pizza?

A: No



Q: Is it sunny outside?

A: Yes



Q: When is the expiration date?

A: unanswerable

Figure 2.1: Examples of different VQA Datasets: (top) CLEVR [86], (middle) VQA [6], (bottom) VizWiz [60]

2.3.2 Methods

Early VQA approaches focused on combining visual features with natural language processing techniques, where visual and textual features were combined using late or early fusion strategies [89, 137]. With the advent of deep learning, VQA witnessed a significant shift towards end-to-end trainable models [138, 245], where CNNs and RNNs were employed to extract visual and textual features respectively. Later on, attention mechanisms methods [5, 96, 184, 233, 249] played a crucial role in improving VQA performance by selectively focusing on relevant image regions or words during the reasoning process. One relevant methods is SAAA [96] which creates multiple attention maps (glimpses) from the image and textual features that are concatenated and passed through an MLP to produce the probabilities over the answer classes. SAN [233] employs multiple layers to progressively attend to relevant image regions and question words, refining the visual and textual representations at each layer. BUTD [5] proposes a combined bottom-up and top-down attention mechanism that enables attention to be calculated at the level of objects and other salient image regions. Instead of performing simple concatenation, sum or dot product between the image and text vectors as most of the methods, MCB [48] uses Bilinear pooling models [202] to compute the outer product between both modality vectors, capturing the relation between all the elements. Finally, in [85] the authors revisit grid features for VQA as an alternative to the popular methods based on BUTD attention, showing that grid-based methods perform an order of magnitude faster while yielding comparable results.

Part I

Natural Scenes

Chapter 3

Scene Text Visual Question Answering

Visual question answering datasets did not consider the rich semantic information conveyed by text within an image. In this chapter, we present a new task aimed to highlight the importance of exploiting high-level semantic information present in images as textual cues in the Visual Question Answering process. We create a dataset ST-VQA, and define a series of tasks of increasing difficulty for which reading the scene text in the context provided by the visual information is necessary to reason and generate an appropriate answer. We also propose a new evaluation metric for these tasks to account both for reasoning errors as well as shortcomings of the text recognition module. In addition, we put forward a series of baseline methods which provide further insight to the newly released dataset, and propose a new model that integrates the textual and visual features by representing the images with a multimodal grid, where a single attention mechanism points to the cell containing the answer.

3.1 Introduction

Textual content in man-made environments conveys important high-level semantic information that is explicit and not available in any other form in the scene. Interpreting written information in man-made environments is essential in order to perform most everyday tasks like making a purchase, using public transportation, finding a place in the city, getting an appointment, or checking whether a store is open or not, to mention just a few. Actually, text is present in about 50% of the images in large-scale datasets such as MS Common Objects in Context [212] and the percentage goes up sharply in urban environments. It is thus fundamental to design models that take advantage of these explicit cues, ensuring that scene text is properly accounted for and is not a marginal research problem, but quite central for holistic scene interpretation models.

The research community on reading systems has made significant progress over the past decade [55, 93], endowing computer vision systems with basic scene text reading capabilities. However, these advancements have not been effectively applied to solve higher-level problems. At the same time, Visual Question Answering (VQA) datasets and models present serious limitations as a result of ignoring scene text content, with disappointing results on questions that require scene text understanding. Therefore, it is timely to bring together these two research lines in the VQA domain to move towards more human like reasoning. However, integrating the textual modality in existing VQA pipelines is not trivial. On one hand, spotting *relevant* textual information in the scene requires performing complex reasoning about positions, colors, objects and semantics, to localize, recognize and eventually interpret the recognized text in the context of the visual content, or any other contextual information available. On the other hand, current VQA models display important dataset biases [86] as well as failures in counting [1, 25], comparing and identifying attributes. These limitations make current models unsuitable to directly integrate scene text information, which is often orthogonal and uncorrelated to the visual statistics of the image.

To this end, we address these limitations by proposing a new task, where the questions and answers are attained in a way that questions can only be answered based on the text present in the image. Accordingly, we create a new dataset called *Scene Text Visual Question Answering* (ST-VQA) which we gather by consciously drawing the majority (85.5%) of the images from datasets that already have generic question/answer pairs. Hence, making possible to combine them with ST-VQA to establish a more generic, holistic VQA task. Some sample images and questions from the collected dataset are shown in Figure 3.1. In addition, we introduce three tasks of increasing difficulty that simulate different degrees of availability of contextual information, and define a new evaluation metric to better discern the models' answering ability employing the Levenshtein distance [114] to account both for reasoning errors as well as shortcomings of the text recognition subsystem [55].

Finally, we propose a novel method based on a single attention mechanism over a grid of multimodal spatial features, where the output weights of the attention module are interpreted as the probability that a certain spatial location of the image contains the answer text to the given question.



Q: What is the price of the bananas per kg?

A: \$11.98



Q: What does the red sign say?

A: Stop



Q: Where is this train going?

A: To New York

A: New York



Q: What is the exit number on the street sign?

A: 2

A: Exit 2



Q: What is the cost for delivery of Espresso Pizza?

A: Free Delivery



Q: What is this dish?

A: Chicken Dijon

Figure 3.1: Recognizing and interpreting textual content is essential for scene understanding. In the Scene Text Visual Question Answering (ST-VQA) dataset leveraging textual information in the image is the only way to solve the QA task.

3.2 ST-VQA Dataset

3.2.1 Dataset creation

In this section we describe the process for collecting images, questions and answers for the ST-VQA dataset, and offer an in-depth analysis of the collected data. Subsequently, we detail the proposed tasks and introduce the evaluation metric.

Images: The ST-VQA dataset comprises 23,038 images sourced from a combination of public datasets that include both scene text understanding datasets as well as generic computer vision ones. In total, we used six different datasets, namely: ICDAR 2013 [94] and ICDAR2015 [93], ImageNet [40], VizWiz [60], IIIT Scene Text Retrieval [149], Visual Genome [104] and COCO-Text [212]. A key benefit of combining images from various datasets is the reduction of dataset bias such as selection, capture and negative set bias which have been shown to exist in popular image datasets [98]. Consequently, the combination of datasets results in a greater variability of questions. To automatically select images to define questions and answers, we use the end-to-end single shot text retrieval architecture YOLO-PHOC [53]. We automatically select all images that contain at least 2 text instances, thus ensuring that the proposed questions contain at least 2 possible options as an answer. The final number of images and questions per dataset can be found in Table 3.1.

Original Dataset	Images	Questions
Coco-text	7,520	10,854
Visual Genome	8,490	11,195
VizWiz	835	1,303
ICDAR	1,088	1,423
ImageNet	3,680	5,165
IIIT-STR	1,425	1,890
Total	23,038	31,791

Table 3.1: Number of images and questions gathered per dataset.

Question and Answers: The ST-VQA dataset comprises 31,791 questions. To gather the questions and answers of our dataset, we used the crowd-sourcing platform Amazon Mechanical Turk (AMT). During the collection of questions and answers, we encouraged workers to come up with closed-ended questions that can be unambiguously answered with text found in the image, prohibiting them to ask yes/no questions or questions that can be answered only based on the visual information.

The process of collecting question and answer pairs consisted of two steps. First, the workers were given an image along with instructions asking them to come up with a question that can be answered using the text found in the image. The workers were asked to write up to three question and answer pairs. Then, as a verification step, we perform a second AMT task that consisted of providing different workers with the image and asking

them to respond to the previously defined question. To avoid ambiguous questions, we checked and corrected if necessary all the questions for which we did not obtain the same answer in both steps. In some cases both answers were deemed correct and accepted to capture the language variability. Therefore, ST-VQA questions have up to two different valid answers.

In total, the ST-VQA dataset comprises 31,791 questions/answers pairs posed over 23,038 images, with separated splits of 19,027 images and 26,308 questions for training, and 2,993 images and 4,163 questions for testing. We present examples of question and answers of our dataset in Figure 3.1.

3.2.2 Analysis and Comparison with TextVQA

TextVQA [191] is a concurrent work to the one presented here. Similarly to ST-VQA, TextVQA proposes an alternative dataset for VQA which requires reading and reasoning about scene text. Additionally, it also introduces LoRRA, a novel architecture that combines a standard VQA model Pythia [190] and an independently trained OCR module [21] with a “copy” mechanism, inspired by pointer networks [59, 213], which allows using OCR recognized words as predicted answers if needed. Both TextVQA and ST-VQA datasets are conceptually similar, although there are important differences in the implementation and design choices. In the case of ST-VQA, the images have been collected from different sources, including scene text understanding datasets, while in the case of TextVQA all images come from a single source, the Open Images [103] dataset. In addition, to select the images to annotate for the ST-VQA, we explicitly required a minimum amount of two text instances to be present, while in TextVQA the images were sampled on a category basis, emphasizing categories like “billboard”, “traffic sign”, or “whiteboard”, that are expected to contain text. In terms of questions, ST-VQA focuses on questions that can be answered unambiguously directly using part of the image text as the answer, while in TextVQA any question requiring reading the image text is allowed.

In Figure 3.2 we provide the length distribution for the gathered questions and answers of the ST-VQA dataset, in comparison to the recently presented TextVQA. It can be observed that the length statistics of the two datasets are closely related. To further explore the statistics of our dataset, Figure 3.3 visualizes how the ST-VQA questions are formed. As it can be appreciated, our questions start with “What, Where, Which, How and Who”. A considerable percentage starts with “What” questions, as expected given the nature of the task. A critical point to realize however, is that the questions are not explicitly asking for specific text that appears in the scene; rather they are formulated in a way that requires to have certain prior world knowledge/experience. For example, some of the “*what*” questions inquire about a brand, website, name, bus number, etc., which require some explicit knowledge about what a brand or website is.

There has been a lot of effort to deal with the language prior inside the datasets [57, 86, 238]. One of the reasons for having language priors in datasets is the uneven distribution of answers in the dataset. In VQA v1 [6], since the dataset is formed from the images of




Figure 3.2: Percentage of questions (top) and answers (bottom) per number of words.




Figure 3.3: Distribution of questions in the ST-VQA train set by their starting 4-grams (ordered from center to outwards). Words with a small contribution are not shown for better visualization.




Figure 3.4: Distribution of answers for different types of questions in the ST-VQA train set. Each color represents a different unique answer.

MSCOCO [124], the answers to the question of “what sport ...” are *tennis* and *baseball* over 50%. Another example is the question “is there ...”, having *yes* as an answer in over 70% of the cases. As can be seen in Figure 3.4, our dataset follows a uniform distribution for the answers except for “sign” and “year”, reducing the risk of language priors while having a big vocabulary for the answers.

To put ST-VQA in perspective, VQA 2.0 [57], the biggest dataset in the community, contains 1.1 million questions out of which only 8k, corresponding to less than 1% of the total questions, requires reading the text in the image. The TextVQA [191] dataset, which comprises 28,408 images paired with 45,336 questions, reported that 39% (18k) of the answers do not contain any of the OCR tokens¹. This might be either due to the type of the questions defined, or due to shortcomings of the employed text recognition engine. In contrast, as a result of the collection procedure followed in ST-VQA, all questions (31,791) can be answered unambiguously directly using the text in the image.

Moreover, the fact that ST-VQA answers are explicitly grounded to the scene text allows us to collect a single answer per question. More specifically, to consider an answer as correct we introduce a soft metric that requires the predicted answer to have a small edit distance to the correct answer (see section 3.2.4), factoring this way in the evaluation procedure the performance of the text recognition sub-system. In the case of TextVQA, 10 answers are collected per question and any answer supported by at least three subjects is considered correct. In order to better understand the effects of our approach compared to collecting multiple responses like in TextVQA, we performed an experiment collecting 10 answers for a random subset of 1,000 ST-VQA questions. Our analysis showed that in 84.1% of the cases there is agreement between the majority of subjects and the original answer. The same metric for TextVQA is 80.3%, confirming that defining a single unambiguous answer results in similarly low ambiguity at evaluation time.

Despite the differences, the two datasets are highly complementary as the image sources used do not intersect with each other, creating an opportunity for transfer learning between the two datasets and maybe combining data for training models with greater generalization capabilities.

¹Presentation of the TextVQA Challenge, CVPR 2019

3.2.3 Tasks

We define 3 novel tasks, suitable for the ST-VQA dataset, namely “strongly contextualized”, “weakly contextualized” and “open vocabulary”. The proposed differentiation of tasks can be interpreted by how humans make use of prior knowledge to argue about their current situation. Such prior knowledge in ST-VQA is provided as a dictionary, different for each task. Our formulation of the tasks is inspired by the previous notions available, and the difficulty per task increases gradually. In the strongly contextualized task, we capture this prior knowledge by creating a dictionary *per image* for the specific scenario depicted. In the weakly contextualized task, we provide a *single dictionary* comprising all the words in the answers of the dataset. Finally, for the open dictionary task, we treat the problem as tabula rasa where no a priori and no external information is available to the model. We show some examples of the tasks in Figure 3.5

For the strongly contextualized task (1), following the standard practice used for end-to-end word spotting [93, 94, 217], we create a dictionary per image that contains the words that appear in the answers defined for questions on that image, along with a series of distractors. The distractors are generated in two ways. On one hand, they comprise instances of scene text as returned by a text recognizer applied on the image. On the other hand, they comprise words obtained by the dynamic lexicon generation model TextTopicNet [54, 158]. Hence, the dictionary for the strongly contextualized task is comprised by the ground truth plus a small set of distractors, which in total adds up to 100 words. In the weakly contextualized task (2), we provide a unique dictionary of 30,000 words for all the datasets’ images, which is formed by collecting all the 22k ground truth words plus 8k distractors generated in the same way as in the previous task. Finally, for the open dictionary task (3) we provide no extra information, thus it is an open-lexicon task.

By proposing the aforementioned tasks, the VQA problem is conceived in a novel manner that has certain advantages. First, it paves the way for research on automatically processing and generating such prior information, and its effect on the model design and performance. Second, it provides an interesting training ground for end-to-end reading systems, where the provided dictionaries can be used to prime text spotting methods.

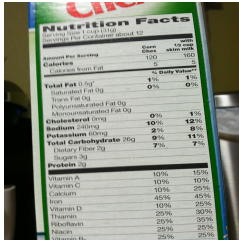


<div></div>			
Question	What does the sign say in the center of the image?	What brand is the bread?	How many calories are from fat?
Answer	Emergency Stop Plunger	An Luc Pele	5
Task 1 Dictionary (40 words out of 100)	HOUSE, NATIONAL, NUMBER, EMERGENCY , SINGLE, STOP , FORCE, STILL, BASED, EVEN, MEMBER, UNITED, GROUP, WEST, BEGAN, DUE, 1, POWER, END, JUNE, FOUND, SLOP, JULY, DAY, POPULATION, KING, JOHN, ISLAND, REGION, SECOND, LARGE, AMERICAN, TEAM, SMALL, SOUTH, PLUNGER , ...	LUC , PLAYER, STILL, 0, HOME, JULY, AIR, OFTEN, WILL, USE, BASED, SERVICE, MUSEUM, SINCE, AN , FRANCE, MEMBER, PELE , POWER, CENTER, PARK, FRENCH, WATER, SUCK, DAY, POPULATION, SCHOOL, SEVERAL, LIKE, ISLAND, REGION, WORK, LARGE, PLACE, AMERICAN, TABLE, ...	SEPTEMBER, YORK, SERIES, NUMBER, GENERAL, HIGH, FACTS, STILL, EVEN, WILL, MEMBER, DEATH, BASED, 2, AREA, MUCH, BEGAN, FRANCE, 0, 3, WON, 5, INTERNATIONAL, JOHN, NORTH, FRENCH, GAME, END, 484, 5 , FOUR, SCHOOL, 787, LAST, LIKE, UNIVERSITY, AIS, LARGE, 533, LEFT, ...
Task 2 Dictionary (100 words out of 30k)	SERVICE, 5 , FACTS, 0, JOHN, REGION, BASED, INTERNATIONAL, WON, FOUR, LARGE, SEVERAL, GENERAL, 533, TEAM, DAY, NORTH, FORCE, POWER, WILL, 0, AMERICAN, NUMBER, SUCK, LIKE, SEPTEMBER, YORK, LAST, WATER, BEGAN, EVEN, UNIVERSITY, DAY, BASED, STILL, LIKE, PLAYER, LARGE, JUNE, PARK, POWER, HOME, FRENCH, BEGAN, RIVER, STOP , MUCH, EMERGENCY , GROUP, STILL, WILL, HOUSE, PLACE, WEST, SINCE, SINGLE, SMALL, ISLAND, OFTEN, FRANCE, JULY, UNITED, SOUTH, LUC , POPULATION, SECOND, REGION, AREA, FOUND, AMERICAN, FOLLOWING, BASED, ISLAND, STILL, CENTER, 3, HIGH, PELE , FRANCE, DEATH, JOHN, EVEN, AIS, LEFT, MEMBER, USE, SLOP, GAME, SERIES, 484, PLUNGER , NUMBER, MUSEUM, SCHOOL, KING, MEMBER, FRENCH, MEMBER, NATIONAL, LARGE, AN , JULY, 1, POPULATION, 787, DUE, WORK, END, SCHOOL, END, AIR, ...		
Task 3 Dictionary	N/A	N/A	N/A
Image Source	ICDAR	Visual Genome	Vizwiz
Reasoning Requirements	Position, Object	Object Relation, A priori	Text Structure

Figure 3.5: Examples of the three tasks with their corresponding provided dictionaries.

3.2.4 Evaluation metric

Since the answers of our dataset are contained within the text found in the image, which is dependent on the accuracy of the OCR being employed, the classical evaluation metric of VQA tasks is not optimum for our dataset, e.g. if the model reasons properly about the answer but makes a mistake of a few characters in the recognition stage, like in Figure 3.8 (first row, third column), the typical accuracy score would be 0. However, the metric we propose named Average Normalized Levenshtein Similarity (ANLS) would give an intermediate score between 0.5 and 1 that will softly penalize the OCR mistakes. Thus, a motivation of defining a metric that captures OCR accuracy as well as model reasoning is evident. To this end, in all 3 tasks we use the Normalized Levenshtein Similarity [114] as an evaluation metric. More formally, we define ANLS as follows:

$$\text{ANLS} = \frac{1}{N} \sum_{i=0}^N \left(\max_j s(a_{ij}, o_{q_i}) \right) \quad (3.1)$$

$$s(a_{ij}, o_{q_i}) = \begin{cases} (1 - NL(a_{ij}, o_{q_i})) & \text{if } NL(a_{ij}, o_{q_i}) < \tau \\ 0 & \text{if } NL(a_{ij}, o_{q_i}) \geq \tau \end{cases}$$

where N is the total number of questions in the dataset, M is the total number of GT answers per question, a_{ij} are the ground truth answers where $i = \{0, \dots, N\}$, and $j = \{0, \dots, M\}$, and o_{q_i} is the network's answer for the i^{th} question q_i . $NL(a_{ij}, o_{q_i})$ is the normalized Levenshtein distance between the strings a_{ij} and o_{q_i} (notice that the normalized Levenshtein distance is a value between 0 and 1). We define a threshold $\tau = 0.5$ that penalizes metrics larger than this value, thus the final score will be 0 if the NL is larger than τ . The intuition behind the threshold is that if an output has an edit distance greater than 0.5 to an answer, meaning getting half of the answer wrong, we reason that the output is the wrong text selected from the options as an answer. Otherwise, the metric has a smooth response that can gracefully capture errors in text recognition. We provide a visual example in Figure 3.6.



Question	What is the number of Boat?		
GT	25688		
	Output	Acc.	ANLS
Method A	25686	100	1.0
Method B	5686	0	0.8
Method C	6	0	0.0

Figure 3.6: Comparison between Accuracy and ANLS evaluation metrics on the results obtained from three different methods applied to a sample question. If the answer is correct (Method A), the method will receive the highest score for both Accuracy and ANLS. However, if a single character is incorrect due to a text recognition error (Method B), Accuracy will be zero despite the method's reasoning being correct. In contrast, ANLS only imposes a slight penalty for recognition errors. However, if the provided answer is very different (Method C), the score will be zero for both accuracy and ANLS, as the Levenshtein similarity will fall below the 0.5 threshold.

3.3 Baselines

The following section describes the baselines employed in this work, as well as an analysis of the results obtained in the experiments conducted. The proposed baselines help us to showcase the difficulty of the proposed dataset and its tasks. Aside from baselines designed to exploit all the information available (visual information, scene text, and the question), we have purposely included baselines that ignore one or more of the available pieces of information in order to establish lower bounds of performance. The following baselines are employed to evaluate the datasets:

Random: As a way of assessing aimless chance, we return a random word from the contextual dictionary provided for each task. For the open dictionary task, we employ the widely used English vocabulary of 90k words proposed by [81].

Scene Text Retrieval: This baseline leverages a single shot CNN architecture YOLO-PHOC [53] that predicts at the same time bounding boxes and a compact text representation (PHOC [3]) of the words in them. It ignores the question and any other visual information of the image. We have defined two approaches: the first (“STR retrieval”) uses the specific task dictionaries as queries to a given image, and the top-1 retrieved word is returned as the answer; the second one (“STR bbox”), follows the intuition that humans tend to formulate questions about the largest text instance in the image. We take the text representation from the biggest bounding box found and then find the nearest neighbor word in the corresponding dictionaries.

Scene Image OCR: The state-of-the-art text recognition model TextSpotter with Explicit Alignment and Attention [67] is used to process the test set images. The detected text is ranked according to the confidence score, and the closest match between the most confident text detection and the provided vocabularies for task 1 and task 2 is used as the answer. In task 3 the most confident text detection is adopted as the answer directly.

Standard VQA models: We evaluate two standard VQA models. The first one, named “Show, Ask, Attend and Answer” [96] (SAAA), consists of a CNN-LSTM architecture. On one hand, a ResNet-152 [65] is used to extract image features with dimension $14 \times 14 \times 2048$, while the question is tokenized and embedded by using a multi-layer LSTM. On top of the combination of image features and the question embedding, multiple attention maps (glimpses) are obtained. The result of the attention glimpses over the image features and the last state of the LSTM is concatenated and fed into two fully connected layers to obtain the distribution of answer probabilities according to the classes.

The second model, named “Stacked Attention Networks” [233] (SAN), uses a pre-trained VGG [186] to obtain image features with shape $14 \times 14 \times 512$. Two question encoding methods are proposed, one that uses an LSTM and another that uses a CNN, both of them yielding similar results according to the evaluated dataset. The encoded question either by a CNN or LSTM is used along with the image features to compute two attention maps, which later are used with the image features to output a classification vector.

Overall, three different experiments are proposed according to the output classification vector. The first, is formed by selecting the most common 1k answer strings in the ST-

VQA training set as in [6]. For the second one, we selected the 5k most common answers so that we can see the effect of a gradual increase of the output vector in the two VQA models. In the third one, all the answers found in the training set are used (19,296) to replicate the wide range vocabulary of scene-text images and to capture all the answers found in the training set.

Fusing Modalities - Standard VQA Models + Scene Text Retrieval: Using the previously described VQA models, the purpose of this is to combine textual features obtained from a scene text retrieval model with existing VQA pipelines. To achieve this, we use the YOLO-PHOC [53] used in STR baselines, which produces a set of text predictions based on pre-defined anchor boxes for each cell of the grid-patched image. Thus, we employ the output tensor before the non-maximal suppression step (NMS) and take the features of the most confident PHOC predictions above a threshold for each grid cell. The selected features form a tensor of size $14 \times 14 \times 609$, which is concatenated with the image features before the attention maps are calculated on both previously described VQA baselines. Afterwards, the attended features are used to output a probability distribution over the classification vector.

3.3.1 Baselines results

The results of all provided baselines according to the defined tasks are summarized in Table 3.2. As a way to compare the proposed Average Normalized Levenshtein Similarity (ANLS) metric, we also calculate the accuracy for each baseline. The accuracy is calculated by counting the exact matches between the model predictions and collected answers, as is the standard practice in the VQA literature.

The last column in Table 3.2, upper bound, shows the maximum possible score that can be achieved depending on the method evaluated. The upper bound accuracy for standard VQA models is the percentage of questions where the correct answer is part of the models' output vocabulary, while the upper bound ANLS is calculated by taking as answer the closest word (output class) in terms of Levenshtein distance to the correct answer. In the case of the Scene Text Retrieval (STR retrieval) [53] model, the upper bound is calculated by assuming that the correct answer is a single word and that this word is retrieved by the model as the top-1 among all the words in the provided vocabularies.

It can be seen that standard VQA models that disregard textual information from the image achieve similar scores, ranging between 0.085 to 0.102 ANLS, or 6.36% to 7.78% accuracy. One relevant point is that although in VQA v1 [6] the SAAA [96] model is known to outperform SAN [233], in our dataset the effect found is the opposite, due to the fact that our dataset and task outline is different in its nature compared to VQA v1. Another important point is that the SAAA model increases both its accuracy and ANLS score when using a larger classification vector size, from 1k to 5k classes. However, going from 5k to 19k classes the results are worse, suggesting that learning such a big vocabulary in a classification manner is not feasible.

Method	OCR	Q	V	Task 1		Task 2		Task 3		Upper bound	
				ANLS	Acc.	ANLS	Acc.	ANLS	Acc.	ANLS	Acc.
Random	✗	✗	✗	0.015	0.96	0.001	0.00	0.00	0.00	-	-
STR [53] (retrieval)	✓	✗	✗	0.171	13.78	0.073	5.55	-	-	0.782	68.84
STR [53] (bbox)	✓	✗	✗	0.130	7.32	0.118	6.89	0.128	7.21	-	-
Scene Image OCR [67]	✓	✗	✗	0.145	8.89	0.132	8.69	0.140	8.60	-	-
SAAA [96] (1k cls)	✗	✓	✓	0.085	6.36	0.085	6.36	0.085	6.36	0.571	31.96
SAAA+STR (1k cls)	✓	✓	✓	0.091	6.66	0.091	6.66	0.091	6.66	0.571	31.96
SAAA [96] (5k cls)	✗	✓	✓	0.087	6.66	0.087	6.66	0.087	6.66	0.740	41.03
SAAA+STR (5k cls)	✓	✓	✓	0.096	7.41	0.096	7.41	0.096	7.41	0.740	41.03
SAAA [96] (19k cls)	✗	✓	✓	0.084	6.13	0.084	6.13	0.084	6.13	0.862	52.31
SAAA+STR (19k cls)	✓	✓	✓	0.087	6.36	0.087	6.36	0.087	6.36	0.862	52.31
SAN(LSTM) [233] (5k cls)	✗	✓	✓	0.102	7.78	0.102	7.78	0.102	7.78	0.740	41.03
SAN(LSTM)+STR (5k cls)	✓	✓	✓	0.136	10.34	0.136	10.34	0.136	10.34	0.740	41.03
SAN(CNN)+STR (5k cls)	✓	✓	✓	0.135	10.46	0.135	10.46	0.135	10.46	0.740	41.03
H1 Human performance [†]	✓	✓	✗	-	-	-	-	0.521	37.54	1.000	100.0
H2 Human performance [†]	✗	✓	✓	-	-	-	-	0.214	18.81	1.000	100.0
H3 Human performance [†]	✓	✓	✓	-	-	-	-	0.851	78.16	1.000	100.0

Table 3.2: Baseline results comparison on the three tasks of ST-VQA dataset. We provide Average Normalized Levenshtein similarity (ANLS) and Accuracy for different methods that leverage OCR, Question (Q) and Visual (V) information. †: Human performance study was performed over a subset of 1,000 questions of the test set.

It is worth noting that the proposed ANLS metric generally tracks accuracy, which indicates broad compatibility between the metrics. But, in addition, ANLS can deal with border cases (i.e. correct intended responses, but slightly wrong recognized text) where accuracy, being a hard metric based on exact matches, cannot. Such border cases are frequent due to errors at the text recognition stage. Examples of such behavior can be seen in the qualitative results shown in Figure 3.8 for some of the answers (indicated in orange color). This also explains why the “Scene Image OCR” model is better ranked in terms of ANLS than of accuracy in Table 3.2.

Standard VQA models that disregard any textual information, perform worse or comparable at best to the “STR (retrieval)” or “Scene Image OCR” models, despite the fact that these heuristic methods do not take into account the question. This observation confirms the necessity of leveraging textual information as a way to improve performance in VQA models. We demonstrate this effect by slightly improving the results of VQA models (SAAA and SAN) by using a combination of visual features and PHOC-based textual features (SAAA+STR and SAN+STR).

For further analysis of the baseline models’ outputs and comparison between them, we provide in Figure 3.7 two bar charts with specific results on different question types. In most of them, the STR model is better than the “Scene Image OCR” (ST-OCR) in terms of ANLS. The effect of PHOC embedding is especially visible on the SAN model for correctly answering the question type such as “what year”, “what company” and “which”. Also, none of the models is capable of answering the questions regarding license plates, “who” and “what number”. This is an inherent limitation of models treating VQA as a




Figure 3.7: Results of baseline methods in the open vocabulary task of ST-VQA by question type.

pure classification problem, as they can not deal with out of vocabulary answers. In this regard, the importance of using PHOC features lies in their ability to capture the morphology of words rather than their semantics as in other text embeddings [20, 147, 159]; since several text instances and answers in the dataset may not have any representation in a pre-trained semantic model. The use of a morphological embedding like PHOC can provide a starting point for datasets that contain text and answers in several languages and out of dictionary words such as license plates, prices, directions, names, etc.


In the human performance study reported in last rows of Table 3.2, we appreciate that when both image, scene text and question are provided to the human evaluator (H3), the performance is consistent with the human study reported in TextVQA [191] in terms of accuracy. Their study shows a human accuracy of 85.0 in TextVQA, but having collected 10 answers per question their accuracy metric is a bit more flexible in accepting diverse correct answers. Moreover, when we blur all text instances in H1, the performance drops immensely, while when we show only the question and the list of OCR words in H2, the performance drop is less significant than in H1. This, demonstrates that the textual cue is much more important than the visual cue in ST-VQA. Another point to stress is that humans are especially good at answering questions without even seeing the image. This is because of the fact that humans use a-priori knowledge such as what a number or a license plate are. As an example, an image for which the question is “*What is the price of ...*” can be correctly answered by selecting a unique numerical OCR token since the price has to be a number.



Q: What brand are the machines?
A: **bongard**
SAN(CNN)+STR: **ray**
SAAA+STR: **ray**
Scene Image OCR: **zbongard**
STR (bbox): **1**



Q: Where is the high court located?
A: **delhi**
SAN(CNN)+STR: **delhi**
SAAA+STR: **delhi**
Scene Image OCR: **high**
STR (bbox): **delhi**



Q: What does the black label say?
A: **GemOro**
SAN(CNN)+STR: **st.**
george ct.
SAAA+STR: **esplanade**
Scene Image OCR: **gemors**
STR (bbox): **genoa**



Q: What's the street name?
A: **place d'armes**
SAN(CNN)+STR: **10th st**
SAAA+STR: **ramistrasse**
Scene Image OCR: **d'armes**
STR (bbox): **dames**




Q: What is the route of the bus?
A: **purple route**
SAN(CNN)+STR: **66**
SAAA+STR: **508**
Scene Image OCR: **1208**
STR (bbox): **purple**



Q: What is the automobile sponsor of the event?
A: **kia**
SAN(CNN)+STR: **kia**
SAAA+STR: **kia**
Scene Image OCR: **kin**
STR (bbox): **0**



Q: Which dessert is showcased?
A: **donut**
A: **Vegan Donut**
SAN(CNN)+STR: **t**
SAAA+STR: **Donuts**
Scene Image OCR: **175**
STR (bbox): **north**



Q: What is preheat oven temperature?
A: **350**
SAN(CNN)+STR: **350**
SAAA+STR: **0**
Scene Image OCR: **high**
STR (bbox): **receivables**

Figure 3.8: Qualitative results for different methods on task 1 (strongly contextualised) of the ST-VQA dataset. For each image we show the question (Q), ground-truth answer (blue), and the answers provided by different methods (green: correct answer, red: incorrect answer, orange: incorrect answer in terms of accuracy but partially correct in terms of ANLS ($0.5 \leq \text{ANLS} < 1$)).

3.4 MMGrid Method

Even though adding textual information to generic VQA models leads to improvements, we also show that ad-hoc baselines can outperform them, reinforcing the need of different approaches. Furthermore, existing VQA models usually address the problem as a classification task, but in the case of scene text based answers, the number of possible classes is intractable. Therefore, we propose a new model illustrated in Figure 3.9, which consists of four different modules: an image encoder (CNN), a scene text encoder (OCR + FastText), a question encoder (LSTM + FastText), and the answer prediction module. Overall, MMGrid builds a grid of multimodal features by concatenating convolutional features and a spatial aware arrangement of word embeddings, so that the resulting grid combines the features of the two modalities at each spatial location (cell). Then it uses an attention module that attends to the multimodal spatial features conditioned to the question. The output weights of this attention module are interpreted as the probability that a certain spatial location (grid cell) of the image contains the answer to the given question. Accordingly, the OCR from the cell with highest probability is chosen as the final answer.

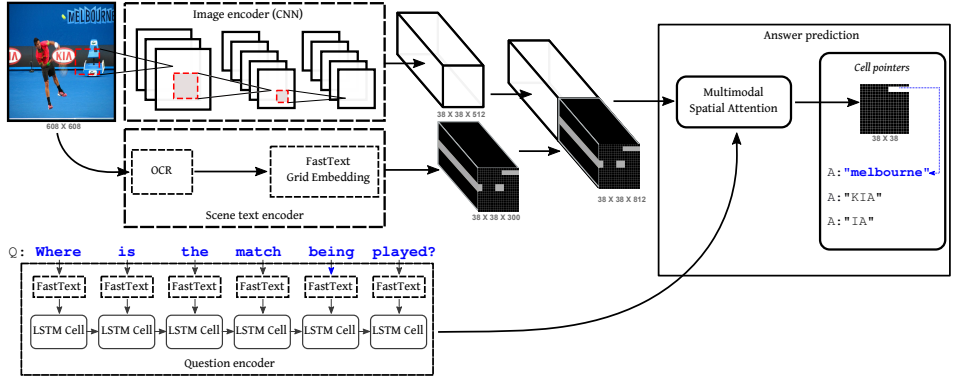


Figure 3.9: Our scene text VQA model consists in four different modules: a visual feature extractor (CNN), a scene text feature extractor (OCR + FastText), a question encoder (LSTM + FastText), and the answer prediction model.

Image encoder: The image encoder is a grid-based visual feature extraction $f_{CNN}(I)$ based on the architecture of the YOLOv3 [169] model with weights pre-trained on the MS-COCO dataset. The YOLOv3 model has a total of 65 successive 3×3 and 1×1 convolutional layers and residual connections. We extract features from the 61st layer, which produces a feature map with dimensions $38 \times 38 \times 512$ that encode high-level object semantics. This configures the features' grid size in our model to be 38×38 . The size of the grid is chosen so that we can quantize the textual information without losing small words (see next section). A 38×38 grid size means each cell corresponds to a 16×16 patch of the input image (with an 608×608 resolution), which means that the smallest possible bounding box of a text instance we expect to find is 16×16 .

Scene text encoder: The first step in our textual feature extractor $f_{ST}(I)$ is to use an optical character recognition (OCR) model to the input image in order to obtain a set of word bounding boxes and their transcriptions $T = \{(b_1, t_1), (b_2, t_2), \dots, (b_n, t_n)\}$. For this, we make use of the Google OCR API [56] for all the experiments. Then, we embed the words extracted from the OCR module into a semantic space by using a pretrained word embedding model. In our case we make use of the FastText word embedding by [20], because it allows us to embed out of vocabulary (OOV) words. Notice that OOV words are quite common in scene text VQA because of two reasons: first, some question may refer to named entities or structured textual information that is not present in closed vocabularies, e.g. telephone numbers, e-mail addresses, website URLs, etc.; second, the transcription outputs of the OCR may be partially wrong, either because the scene text is almost illegible, partially occluded or out of the frame. We use the FastText pretrained model with 1 million 300d word vectors, trained with subword information on Wikipedia 2017, UMBC webbase corpus and statmt.org news dataset, and not updated during training.

With all word transcriptions in T embedded in the FastText 300d space we construct a $38 \times 38 \times 300$ tensor by assigning each of their bounding boxes to the cells in a 38×38 grid with which they overlap as illustrated in Figure 3.10, so that the embedding vectors maintain the same relative spatial positions as the words in the original image. In order to overcome small words being overlapped by larger words, we do this assignment in order, from larger to smaller words. The cells without any textual information are set to zero value. Finally, we concatenate the outputs of the image encoder and the scene text encoder to obtain the multimodal grid based features of the image $f_m(I) = [f_{CNN}(I); f_{ST}(I)] \in \mathbb{R}^{38 \times 38 \times 812}$.




Figure 3.10: Grid cell assignment of the OCR words' bounding boxes. Given an input image (a), the bounding boxes of the words extracted from the OCR model (b) are assigned to their overlapping cells.

Question encoder: The question encoder is based on a LSTM [52] encoder. Given a question Q with N words $Q = \{q_1, q_2, \dots, q_N\}$ we first embed each word with the FastText word embedding, and then feed each word embedding vector into the LSTM. The final hidden layer of the LSTM model is a vector with 1024 dimension and is taken as the output of the question encoder.

Answer prediction module: The main component of the answer prediction is an attention mechanism that attends to the spatial multimodal features $f_m(I)$ conditioned on the question embedding $f_q(Q)$. First, the multimodal grid features $f_m(I)$ are convolved by two 1×1 convolutional layers with 1024 and 512 kernels respectively, resulting in a $38 \times 38 \times 512$ tensor. The question encoded vector $f_q(Q)$ goes through a dense layer with 512 output neurons and is tiled/broadcasted to a shape of $38 \times 38 \times 512$. These two tensors (m_{att} and q_{att}) are added and activated with an hyperbolic tangent (tanh) activation. Finally, the resulting tensor of this operation is convolved with a 1×1 convolutional layer with a sigmoid activation function to produce the output attention map p_{att} with shape $38 \times 38 \times 1$:

$$p_{att} = f_{Att}([f_{CNN}(I); f_{ST}(I)], f_q(Q)) \quad (3.2)$$

At this point, we interpret the values in the output attention map p_{att} as the probability of each image cell to contain the correct answer to the given question Q . Notice that by applying a sigmoid activation function to the last convolution layer, we treat the probability for each cell as an individual binary classification problem. This is intentional, as in most of the cases, the bounding box of the correct answer will cover more than one cell. We train our model using the binary cross entropy loss function:

$$E = - \sum_{i=1}^{38} \sum_{j=1}^{38} [g_{i,j} \log p_{i,j} + (1 - g_{i,j}) \log(1 - p_{i,j})] \quad (3.3)$$

where $p_{i,j}$ is the probability value of the cell on the i th row and j th column on the output attention map p_{att} , and $g_{i,j}$ is the ground truth value for that cell: 1 if the cell contains the answer, 0 otherwise. At inference time, the predicted answer is the OCR token assigned to the cell with maximum probability.

We adopt the stacked design of SAN [233] and empirically found an improvement over using a single attention layer. For this we stack two attention layers, in the first one we combine the weighted average over the multimodal spatial features (using the output probability map as weights) with the question embedding (by addition), and this combination is fed to the second attention layer as the question embedding.

Moreover, since our model is made fully convolutional (including the image encoder) on all the visual branch, we can perform inference at different input scales using the same learned weights.

3.5 MMGrid results

We appreciate in Table 3.3 that MMGrid clearly outperforms all previously published methods both in ANLS and accuracy, improving more than 10% ANLS compared to the ST-VQA competition winner VTA (BUTD [5] based method with scene text features), and 5% ANLS over LoRRA. It is important also to recall here that our model is $5\times$ faster than LoRRA at processing an image, as a consequence of using YOLOv3 instead of Faster-RCNN for feature extraction. Notice that Faster R-CNN requires the region-wise inference (second stage) to be computed for the top-300 proposals provided by the Region Proposal Network (first stage), while in YOLO a single forward pass is computed for the whole image.

Table 3.3 also shows the performance comparison on the validation set of TextVQA. In this case we also compare the accuracy in the specific subset of questions for which the answer is among the OCR tokens (indicated as Acc^\dagger in the table), to understand how the presence of answers that do not correspond to scene text instances in the image (e.g. “yes”/“no” answers) affect the performance of our model. In this subset, our model outperforms previous state of the art by a clear margin, while in the whole validation set we observe the opposite. Notice that this is expected because our model has no mechanism

Method	ST-VQA		TextVQA	
	ANLS	Acc	Acc [†]	Acc
SAAA [96]	0.087	6.66	9.09	13.33
SAN [233]	0.102	7.78	–	–
SAN+STR	0.136	10.34	–	–
QAQ [17]	0.256	19.19	–	–
VTA [17]	0.282	18.22	–	–
LoRRA [191]	0.331	21.28	32.03	27.48
MMGrid	0.381	26.06	37.60	21.88
MMGrid + SAAA	–	–	37.60	26.07

Table 3.3: MMGrid performance on ST-VQA test set and TextVQA validation set. Numbers with [†] are from the official implementation of LoRRA trained on ST-VQA using the same OCR tokens as in our model. Acc[†] refers to the subset of questions with answers among OCR tokens.

for providing valid answers to questions for which their answers are not among the OCR tokens, while LoRRA model can cope with these questions by using a fixed vocabulary answer output space similar to standard VQA models. In order to provide a fair comparison in the whole validation set of TextVQA we have combined the predictions of our model with the well known standard VQA model SAAA [96]. In this experiment we have trained the SAAA model on TextVQA with a fixed output space of the most common 3,000 answers, and the results of entry MMGrid+SAAA correspond to an ensemble model in which the answer is selected with a threshold-based decision. More specifically, the ensemble selects the SAAA answer if its classification confidence is above a given threshold, otherwise it selects the answer of our model. We use a threshold decision over the classification score of the SAAA model and not over ours because we have experimentally found that the confidences of SAAA are better indicators for whether a given question can be answered or not without reading the scene text. The threshold value used was set to 0.5 as in a binary classification problem. We appreciate that this ensemble model achieves competitive performance to the state of the art. While SAAA alone has a marginal performance in TextVQA, the confidences of its predictions are good indicators for whether a given question can be answered without reading the scene text. In such a scenario, a model like ours can be leveraged in a mixed dataset where questions may or may not require answers from the OCR tokens’ set.

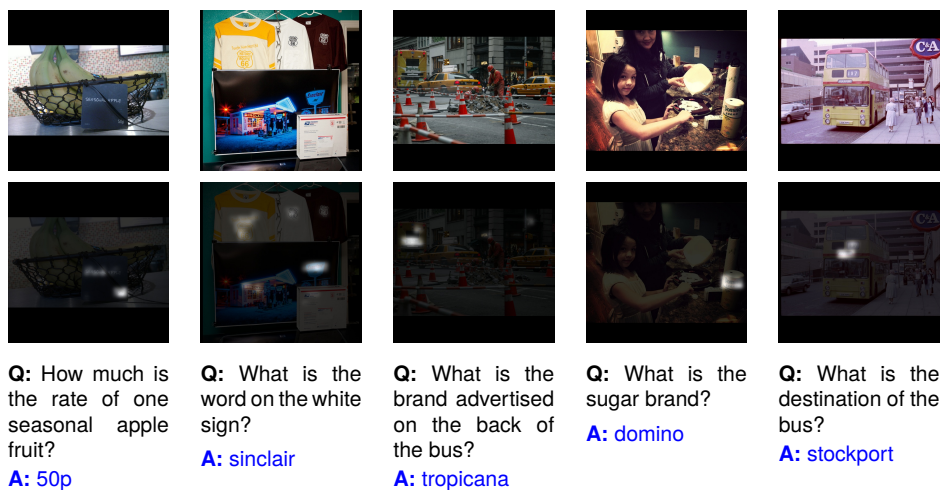


Figure 3.11: Questions correctly answered by MMGrid with the produced attention masks.

Figure 3.11 shows qualitative examples of the produced attention masks and predicted answers for 5 image/question pairs from the ST-VQA test set that are correctly answered by our model. Among them, we can see examples in which textual information alone would suffice to provide a correct answer, but also cases where a joint interpretation of visual and textual cues is needed.

3.6 Evolution of ST-VQA

Dataset	Questions	Images	Image sources	Publication date
TextVQA [191]	45.3 K	28.4 K	Generic CV [103]	CVPR 2019
ST-VQA	31.8 K	23.0 K	ST Understanding [93, 94, 149, 212] Generic CV [40, 60, 104]	ICCV 2019
OCR-VQA [150]	1 M	207 K	Book covers [80]	ICDAR 2019
Text-KVQA [187]	1.3 M	257K	Book covers [80] Movie posters [154] Brand Scene Images	ICCV 2019
STE-VQA [221]	28.0 K	25.2K	English ST [35, 93, 94, 129, 153, 212] Chinese ST [194]	CVPR 2020

Table 3.4: Scene Text VQA datasets

After the publication of ST-VQA and TextVQA works and datasets, the field received a lot of attention. On one hand, several datasets tackled Scene-Text VQA problems while introducing extra challenges and tasks. OCR-VQA [150] is very similar in nature, it is much bigger, but the questions and answers are created automatically from template questions. It consists of more than 1 million question-answer pairs about 207K images from the Book Covers dataset [80], and consequently the questions in this dataset always refer to information that can be found in the book covers such as the authors, title or genre among others. However, in OCR-VQA more than 50% of the questions have answers that are not scene text instances, including 40% binary (yes/no) questions and 10% questions about book genres. Then, Text Knowledge-enabled VQA (Text-KVQA) [187] extended this dataset to enable knowledge based questions by creating 1.3 million question-answer pairs over 257K images and associated web-scale knowledge bases. Unlike previously defined datasets, some of the questions in the Text-KVQA dataset may not be directly answerable from the visual and textual information in the image, and require some reasoning using the provided knowledge bases. The images in this dataset were sourced from Movie Posters [154], Book Covers [80] and a collection of scene images of popular business brands collected by the authors using the Google image search engine. Finally, the Scene Text+Evidence Visual Question Answering (STE-VQA) [221] introduces cross-language and localization challenges in Scene Text VQA. It is a bilingual dataset English-Chinese that comprises 28,062 question/answer pairs on 25,239 images collected from publicly available English and Chinese scene text detection and recognition datasets. Moreover, the question-answer pairs were annotated with their corresponding bounding boxes in the image, which allows models to point to the answer evidences.

Regarding methods, the ST-VQA task has become very popular. Since the publication of the dataset, we have received 1,963 submissions from 73 different teams, out of which 153 methods have been made public. In this review of the task’s evolution, we will focus on the Open Dictionary task, which even though it’s the most challenging, it is the most realistic and has received most of the attention




Figure 3.12: Evolution of method’s performance in ST-VQA dataset. The dashed red line shows the human performance

The method VTA integrated textual features in BUTD [5] and become the winner of the competition held in ICDAR 2019. Afterwards, we proposed the described method MMGrid. One of the majors breaks was brought by the Multimodal Multi-Copy Mesh (M4C) [71], which was the first method to use Transformers [211], fusing all different modalities (question, scene text and visual features) after projecting them into the same dimensional space and concatenating. It also used multiple representations for each of the modalities to cover the different required information to answer the questions, for instance, the recognized words were represented with FastText [20], PHOC [3], the bounding box position and the visual features within the box, each projected into the same dimensional space and averaged. The other major contribution of M4C was the use of an iterative decoder with a pointer network that allowed the model to select a word from a fixed vocabulary or from the recognized words at each time-step.

Most of the subsequent works [49, 61, 92] focused on directing the attention solely to the essential tokens by using Graph Neural Networks or dynamic mask within Transformers, usually based on the position of the tokens. The next big step was brought by Text-Aware Pre-training (TAP) [230] model, which introduced pretraining in ST-VQA task to align the different modalities. Even tough TAP was based on M4C architecture with minor improvements, the pretraining task significantly boosted its performance. Consequently, the next works [99, 134, 215] developed more sophisticated pretraining tasks while still directing the attention towards the desired regions and tokens.

Finally, the rise of multitask multimodal LLM [28, 29, 156, 214] pretrained on vast amounts of data and with several billion parameters has led to results very close to Human Performance.

3.7 Conclusions and Future Work

In this work we introduced a new and relevant dimension to the VQA domain. We proposed the Scene Text VQA, a new dataset for Visual Question Answering that aims to highlight the importance of properly exploiting the high-level semantic information present in images in the form of scene text to inform the VQA process. The dataset comprises questions and answers of high variability, and poses extremely difficult challenges for current VQA methods. We thoroughly analyzed the ST-VQA dataset performing a series of experiments with baseline methods, which established the lower performance bounds, and provided important insights. Although we demonstrate that adding textual information to generic VQA models leads to improvements, we also show that ad-hoc baselines (e.g. OCR-based, which do exploit the contextual words) can outperform them, reinforcing the need for different approaches. Moreover, existing VQA models usually address the problem as a classification task, but in the case of scene text based answers, the number of possible classes is intractable.

To address these limitations, we presented a new model based on an attention mechanism that attends to multimodal grid features, allowing it to reason jointly about the textual and visual modalities in the scene. The provided experiments demonstrate that attending to multimodal features is better than attending separately to each modality. Our grid design choice has also proved to work well in this task, as well as the choice of a one-stage object detection backbone instead of a classification one. Moreover, we have shown that the proposed model is flexible enough to be combined with a standard VQA model, obtaining state-of-the-art results on mixed datasets with questions that can not be answered directly using OCR tokens. However, the method is limited to provide short answers and a generative pipeline such as the ones used in image captioning is required to capture multiple-word answers.

The proposed Average Normalized Levenshtein Similarity metric is better suited for generative models compared to evaluating classification performance, while at the same time, it has a smooth response to the text recognition performance.

After the presentation of this work that jointly with TextVQA pioneered the research field of VQA integrating recognized text from the images, many other works have used the dataset and continued to improve the performance on this task. While the trend has become to use Transformers with different pre-training tasks to align the different modalities. Some methods have also worked on focusing the self-attention to avoid learning redundant or unnecessary features. However, the LLMs have been the ones who have managed to achieve performance close to human.

Part II

Documents

Chapter 4

Background

In this section, we provide an introductory overview of Document Understanding and Machine Reading Comprehension research fields. In the former section we contextualize the existing document tasks that serve as the foundation for the Document Visual Question Answering (DocVQA) task, contextualizing the subsequent chapters 5, 6, 7, 8. In the Machine Reading Comprehension (MRC) related work section, we show previous datasets and methods that perform Question Answering on machine-readable text, which have later inspired several methods proposed on DocVQA.

4.1 Document Understanding

Document understanding has been largely investigated within the document analysis community with the final goal of automatically extracting relevant information from documents, which comprises a varied set of tasks such as character recognition [43], signature verification [79], table extraction [95] or key-value pair extraction [157] among others. And, most of the works have focused on structured or semi-structured documents such as forms, invoices, receipts, passports, ID cards, e-mails or contracts. While earlier works [41, 176] were based on a predefined set of heuristic rules that required the definition of specific templates for each new type of document, learning-based methods [36, 157] allowed to automatically classify the type of document and identify relevant fields of information without predefined templates. Recent advances in deep learning [126, 228, 239] leverage natural language processing, visual feature extraction and graph-based representations to achieve a holistic document understanding that integrates word semantics with the visual layout. The extracted information is then fused using Graph Convolutional Networks (GCN) [102] or Transformer [211] networks to output the required document information. Moreover, recent trends [228] involve pre-training the methods on task-agnostic objectives such as masking and reconstructing tokens from the recognized text, or document classification. These objectives effectively align the in-

terplay between the different input modalities, including semantic meaning, 2-D position, and visual information. Leveraging these pre-training objectives often requires fewer annotations, enabling the methods to be trained on significantly larger datasets than those available for downstream tasks, which leads to better generalization capabilities and performance.

However, all these methods specifically designed for different tasks mainly focus on extracting key-value pairs or finding specific isolated words by following a bottom-up approach from the document features to the relevant semantic information. In contrast, by using the visual question answering paradigm, the different information extraction tasks can be tackled by a single model, following a top-down approach, where the question will guide the information extraction process to find the answer.

Another common application on documents is document retrieval, which aims to search for relevant information in large collections of documents. Most works have addressed this task from the perspective of word spotting [168], i.e., searching for specific query words in the document collection without relying on explicit noisy OCR. Current state-of-the-art on word spotting is based on similarity search in a common embedding space [3] where both the query string and word images can be projected using deep networks [105, 193]. In order to search for the whole collection, these representations are combined with standard deep learning architectures for object detection in order to find all instances of a given word in the document [105, 222]. However, word spotting only allows searching for the specific instances of a given word in the collection, without taking into account the semantic context where that word appears. Instead, when framing a question over the whole collection of documents as we introduce in Chapter 7, the retrieval is not limited to find specific isolated words, but to make a semantic retrieval of documents based on the queried question.

4.2 Machine Reading Comprehension and Question Answering

In the natural language processing community, the use of text in computer readable format is prominent in machine reading comprehension (MRC) and open-domain question answering (QA). In MRC, the goal is to answer a question based on a given passage of text, and it is often formulated as an extractive QA problem where the answer is defined as a span of the context where the question is defined. Some of the popular datasets used for MRC are SQuAD [167], NewsQA [209] and TriviaQA [88], which define questions on Wikipedia, CNN news articles and the Internet respectively. In contrast, MS MARCO [155] proposes an abstractive QA dataset, where the answer is not always present explicitly in the passage, which requires the models to have a deeper understanding of the question and the context to infer the answer. Consequently, abstractive QA is considered more challenging than extractive QA since it involves generating a new text to answer the question rather than selecting a span of text.

In Open-domain QA, the machine is required to find the answer without context, or from a very large collection of unstructured documents or knowledge bases. SearchQA [45] is a dataset created from real queries in Google search engine, and the answers must be retrieved from the search results. Natural Questions (NQ) [107] follows a similar, but it includes two versions of the answers for each question, a long and a short one. In WebQuestions [196], the method needs to gather for extra information in a knowledge base to answer the question. In contrast, HotpotQA [232] is a dataset designed to test systems' ability to answer questions that require both context understanding and multi-hop reasoning. The dataset consists of questions with supporting documents, and the answers often require aggregating information from multiple passages.

Closely related to the topic of this thesis, FigureQA [91] is a synthetically generated dataset over figures, which questions ask for different measures (maximum, minimum, median, area-under-the-curve, ...) on different types of plots. Even though it is required to read the text of the figures to answer the questions, all of them are binary answered yes/no, and the labels used for the plots in different figures are fixed. In this regard, DVQA [90] is a concurrent and similar work, which contains more complex questions divided into three different types: Structure Understanding, Data Retrieval, and Reasoning, which allows assessing different kind of reasoning abilities. However, the text in both datasets is rendered with good quality and standard font types, representing a less challenging setup. Moreover, both datasets lack of the context of a whole document page, and the methods can focus on solving the task without considering text-dense images, tables or complex layouts that can be found in documents. Another similar work, the Textbook Question Answering (TQA) dataset [97] proposed multimodal questions given a context of text, diagrams and images, but the textual information is provided in computer readable format.

In traditional MRC and QA tasks, the context is typically provided in a machine-readable format, and therefore, the main difference between these tasks and VQA relies on the fact that they don't require any kind of visual cue to answer the questions while in DocVQA, the context consists of document images with a variety of visual and textual elements where the models will need to extract the text, and combine computer vision and natural language processing techniques to accurately integrate both modalities to answer the questions. Nevertheless, some datasets for MRC are effectively used for pretraining methods that are afterwards fine-tuned in DocVQA tasks. Moreover, language-only methods provide valuable insights as visually-blind baselines, showing the capabilities of such methods, and the need to integrate the images in multimodal pipelines.

Regarding methods, the earliest methods for Extractive QA [167] combined statistical approaches such as TF-IDF with Part of Speech tagging features and sliding window to retrieve the answer from the given context. With the advent of Neural Networks, word embeddings [20, 147, 159] combined with Recurrent Neural Networks [33, 69, 177] had a significant impact on both extractive and abstractive question answering methods [27, 78, 113, 145, 179], enabling them to capture semantic relationships and contextual information more effectively.

Recent Transformer based pretraining methods like BERT [42] and XLNet [231] have allowed QA models to outperform Humans on reading comprehension like SQuAD [167]. These methods concatenate the question and the context with a special separator token [SEP], and use a segment embedding to distinguish between the question and the passage. Along with the semantic representation, they also use a 1-D positional encoding to provide information about the position of each token in the sequence. The input is then fed through multiple layers of Transformer layers, where each layer applies self-attention to capture the relationships between words in a contextual manner. The self-attention mechanism allows the transformer-based models considering the entire input sequence when generating representations for each token. The output consists of the probability distribution over each token, indicating the likelihood of it being the start or end position of the answer span.

Chapter 5

Visual Question Answering on Single Page Documents

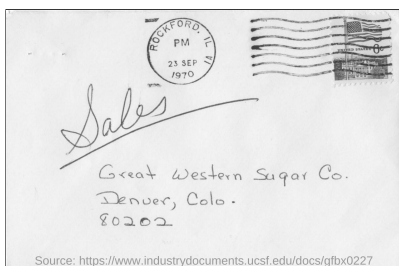
We present a new dataset for Visual Question Answering (VQA) on document images, named SP-DocVQA, which includes more than 12,000 document images and 50,000 associated questions. We also provide a detailed comparative analysis of the dataset with other similar VQA and reading comprehension datasets, and evaluate several baselines based on existing VQA and reading comprehension models. Although these models perform reasonably well on certain types of questions, they show a large performance gap compared to human performance (94.36% accuracy), especially on questions where understanding the structure of the document is crucial.

5.1 Introduction

Research in Document Analysis and Recognition (DAR) has been mainly focused on automatic information extraction from document images through tasks such as optical character recognition [43], table extraction [95] or key-value pair extraction [157]. As described in section 4.1, such algorithms tend to be designed as task specific blocks, blind to the end-purpose the extracted information will be used for. While these information extraction pipelines have progressed independently, achieving holistic document image understanding through a simple constructionist approach built upon such modules is not possible.

In this chapter, we introduce Single Page Document Visual Question Answering (SP-DocVQA), as a high-level task dynamically driving the methods to conditionally interpret the document images based on questions, fostering a “purpose-driven” approach in DIAR research. In SP-DocVQA, as illustrated in Figure 5.1, an intelligent reading system is expected to respond to ad-hoc requests for information, expressed in natural language questions by human users. To do so, reading systems should not only extract and interpret the textual (handwritten, typewritten or printed) content of the document images, but exploit other numerous visual cues including layout (page structure, forms, tables), non-textual elements (marks, tick boxes, separators, diagrams) and style (font, colors, highlighted text), to mention just a few.

Therefore, generic VQA [57] and Scene Text VQA [18, 191] methods are not suitable for document images. Document VQA requires a different approach to leverage all the aforementioned visual cues, but also make use of prior knowledge implicit in written communication conventions, handles the dense semantic information within such images. Furthermore, like ST-VQA, Document VQA answers cannot be sourced from a closed dictionary, since they are inherently open-ended. In contrast to previous approaches that focused on specific document elements such as data visualizations [90, 91] or specific collections such as book covers [150], which are often based on synthetic data and/or template questions lacking diversity. The SP-DocVQA dataset is a large scale, varied collection of real documents, comprising a multitude of different document types including letters, forms, reports or contracts among other, and a range of different textual, graphical and structural elements like tables, forms, or figures, with questions defined by human annotators.



Source: <https://www.industrydocuments.ucsf.edu/docs/gfbx0227>

Q: What is the ZIP code written?

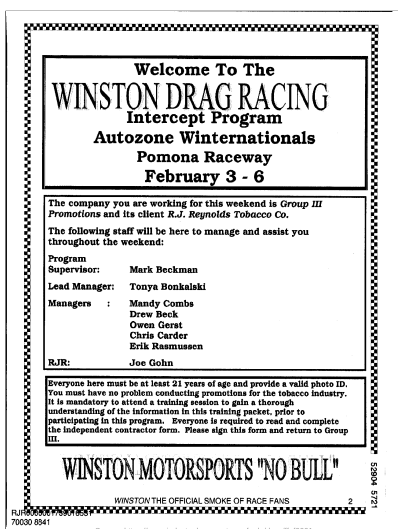
A: 80202

Q: What date is seen on the seal at the top of the letter?

A: 23 sep 1970

Q: Which company address is mentioned on the letter?

A: Great western sugar Co.



Q: What is the flyer about?

A: Winston Drag Racing Intercept Program

Q: When is the Intercept program scheduled?

A: February 3 - 6

Q: Who is the lead manager?

A: Tonya Bonkalski

OUT TOBACCO: MT-768 D BBT391				BEST PROTOTYPE. 327391			
BLERO: BW-6071							
		LBS AT 12.5 %	SOLID LBS			LBS AT TARGET	
STRIPS:	FLUE CURED.....	3,681.7	3,221.5			3,790.0 @ 15.0 %	
	RECONSTITUTED.....	1,996.3	(1,746.8)				
	CASINGS.....	2,159.0	2,159.0			2,540.0 @ 15.0 %	
	ORIENTAL.....	1,243.4	1,089.0			1,280.0 @ 15.0 %	
	RECONSTITUTED.....	2,321.7	2,031.5			2,390.0 @ 15.0 %	
TOTAL STRIPS.....		9,243.1	8,500.0			10,000.0 @ 15.0 %	

GRADE BIN	GRADE	LBS AT 12.5 %	SOLID LBS	LBS AT TARGET
1 C2F	C2F.....	913.1	799.0	945.0 @ 15.0 %
3 L2F	L2F.....	1,311.4	1,147.5	1,350.0 @ 15.0 %
5 T2F	T2F.....	582.9	510.0	600.0 @ 15.0 %
7 L2F	L2F.....	583.2	510.3	600.0 @ 15.0 %
	T2F.....	251.1	254.7	259.7
		874.3	765.0	800.0
65 KBR	KBR.....	1,350.3	1,181.5	1,390.0 @ 15.0 %
70 CPCLF	CPCL.....	971.4	850.0	1,000.0 @ 15.0 %
73 ADR	CAB.....	207.6	181.7	213.8 @ 15.0 %
	YAB.....	444.1	392.3	405.2
		207.6	181.7	213.8
		444.1	392.3	405.2
				1,280.0
977 BEST	MT-778	1,996.6	(1,746.8)	
	+ CASINGS.....		2,159.0	2,540.0 @ 15.0 %

553232685

Source: <https://www.industrydocuments.ucsf.edu/docs/gfbx0227>

Q: What is the Best Prototype mentioned?


A: 327391

Q: What is the LBS at target of total strips?

A: 10,000.0 @ 15.0%

Q: What is the solid LBS of T2F grade?

A: 510.0



Q: When was the MSA segment share dated?

A: October 1995

Q: What is the % of PM?

A: [49.45%, 49.45]

Q: How much is B&W?

A: [19.66%, 19.66]

Figure 5.1: Example question-answer pairs from SP-DocVQA. Answering questions in the new dataset require models not only reading the text, but also interpret it within the layout and structure of the document.

5.2 SP-DocVQA Dataset

5.2.1 Dataset creation

Document Images: Images in the dataset are sourced from documents in the UCSF Industry Documents Library¹. The library is a digital archive of documents from various industries concerning public health and currently hosts more than 40 million documents. We downloaded the documents minimizing binarized images to avoid having poor image quality, and prioritized pages with tables, forms, lists and figures over pages which only contained running text.

The final set of images in the dataset is drawn from 12,767 pages of 6,071 industry documents. More specifically, the documents belong to the 5 major industries the UCSF-IDL library hosts: Tobacco, food, drug, fossil fuel and chemical. However, as observed in Figure 5.3, which shows the industry wise distribution of the 6071 documents, most of them belong to food and nutrition related collections, as their document images are usually non-binarized, unlike in the other industries. Moreover, the documents are dated from as early as 1900 to as recent as 2018 (Figure 5.2a) although the majority are dated from the 1960-2000 period. The document images include typewritten, printed, handwritten and born-digital text, and they represent a wide variety of types including letters, emails, forms, articles, invoices or budgets among others, as it is shown in Figure 5.2b.

Questions and Answers: Questions and answers on the selected document images were collected with the help of remote workers using a Web based annotation tool. To ensure good quality, the workers for the process were selected from a pool of applicants, based on an online quiz assessing applicants' ability to frame questions, adhering to the annotation standards and policies instructed in a prior web tutorial. Moreover, during the annotation process, we provide feedback to the annotators to keep improving annotation quality.

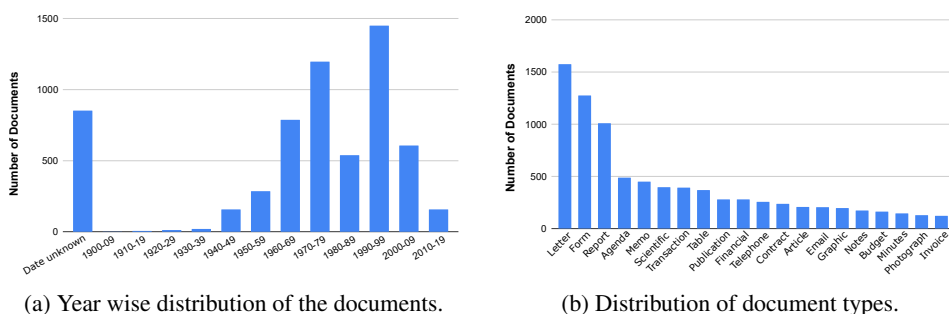


Figure 5.2: Document images are sourced from 6071 documents spanning many decades, of a variety of types, originated from 5 different industries.

¹<https://www.industrydocuments.ucsf.edu/>

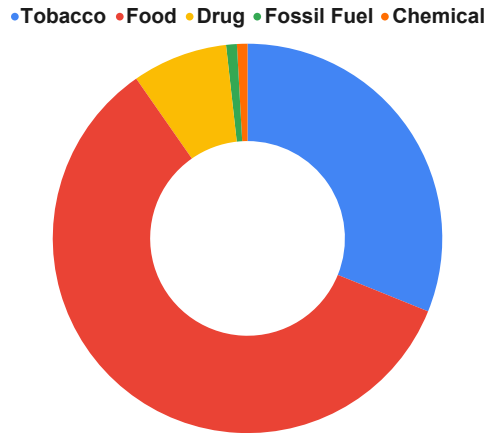


Figure 5.3: Industry-wise distribution of the documents.

The annotation process was organized in three stages, similarly as in ST-VQA. In stage 1, the workers were given a document image and asked to define at most 10 question-answer pairs on it. We encouraged the workers to add more than one ground truth answer per question in cases where it was required to capture language variability. Some examples of these cases are the second and third questions of the bottom-right image in Figure 5.1, for the second question, both 49.5% or just 49.5 are possible correct answers to the question asking about the percentage. The workers were also instructed to ask questions which can be answered using text present in the image, and to enter the answer verbatim from the document. Thus, in the previous example, “forty-nine point five percent” is not considered a possible answer. Even though this makes SP-DocVQA dataset an extractive QA problem, it helps to keep control over the annotations and their quality. Nonetheless, the questions can still be posed on figures or images, as long as the answer’s text is present in the image.

The second annotation stage was aimed to verify the data collected in the first stage. Here the workers were shown an image and the questions defined on it during the first stage (but not the answers), and were required to enter the answers for the questions. In this stage, they were also required to assign one or more question type to each question, which indicates the information or evidence required to answer the question. The different question types in SP-DocVQA are further discussed in statistics section 5.2.2. During the second stage, the workers had the option to indicate that a question was inadequate either due to language issues or ambiguity, and consequently, such questions were not included in the dataset.

If none of the answers entered during the first stage match exactly with any of the answers from the second stage, the question was sent for review in a third stage. Here questions and answers are editable and the reviewer either accepts the question-answer (after editing if necessary) or rejects it. The third stage review was done by the authors themselves.

5.2.2 Statistics and Analysis

The SP-DocVQA dataset comprises 50,000 questions framed on 12,767 images, and the data is split randomly in a 80% – 10% – 10% ratio to train, validation and test splits.

Question types: Figure 5.4 shows the frequency of the 9 question types in the dataset. The questions type(s) collected during the second stage of the annotation process indicates the type of data on which the question is grounded, i.e. the kind of information that needs to be analyzed to answer the question. For example, ‘table/list’ is assigned if answering the question requires to analyze and understand tabular or list structures, like in the second and third questions of the top-right image of Figure 5.1. If the information is in the form of a key:value, the ‘form’ type is assigned like the first question of the same example. ‘Layout’ is assigned for questions that require spatial/layout information to find the answer, like questions asking for a title or heading, which require understanding the structure of the document. If the answer to a question is based on information in sentences or paragraphs, the type assigned is ‘running text’. When the question asks about information on a figure or diagram like in Figure 5.1 (bottom-right), the ‘figure’ label is assigned. Instead, ‘photograph’ type is assigned when the question is posed over a picture taken with a camera like in Figure 5.11 (right). Moreover, for all the questions where the answer is based on handwritten text, the ‘handwritten’ type is assigned. Note that each question can have more than one type associated with it. We provide examples of SP-DocVQA document images with their questions, answers and their types with the qualitative examples in section 5.4.2.




Figure 5.4: The 9 question types and share of questions in each type.

Questions: Figure 5.5a shows the question length distribution in SP-DocVQA compared to other similar datasets. The average question length is 8.12 words, the second highest among the compared datasets, surpassed only by SQuAD (10.07). Additionally, the dataset stands out with 70.72% of its questions being unique, indicating a diverse range of question types and topics. Figure 5.5b shows the top 10 most frequent questions and their frequencies. There are multiple questions asking for dates, titles and page numbers. A sunburst of the first 4 words of the questions is shown in Figure 5.5c, which shows that a large majority of questions start with “what is the”, asking for the aforementioned dates and titles, but also names or total and amount quantities.

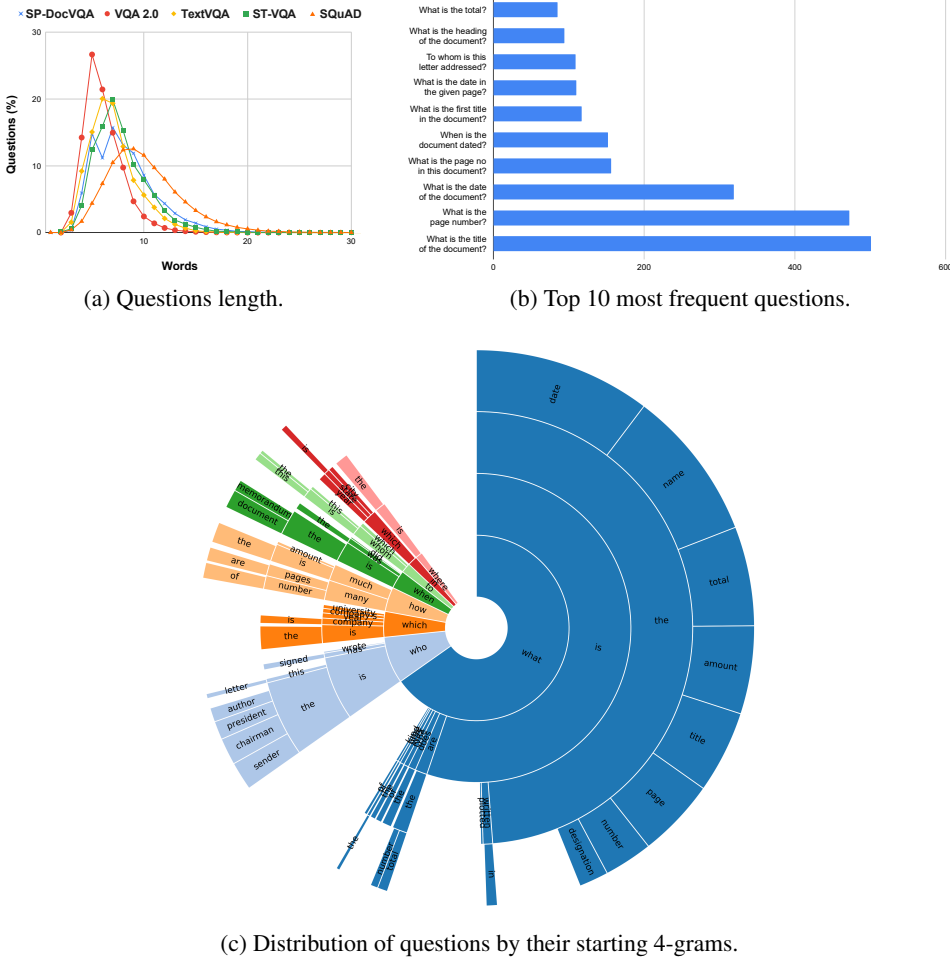


Figure 5.5: Question statistics and comparison with similar VQA and MRC datasets. Most of the questions aim to retrieve relevant information in documents such as date, title, total amount or page number.

Answers: The distribution of answer lengths is shown in Figure 5.6a. We observe in the figure that both SP-DocVQA and SQuAD have a higher number of longer answers compared to the VQA datasets. The average answer length is 2.17, with 63.2% of the answers being unique. The top 10 answers in the dataset are shown in Figure 5.6c. While the answer space is quite diverse (the most common answer “2” is the answer for only 0.7% of questions in the dataset). We observe that almost all the top answers are numeric values, which is expected because there are many document images of reports and invoices. In Figure 5.6d we show the top 10 non-numeric answers, which are mostly named entities about people, institutions and places. This, is strongly correlated with the word cloud in Figure 5.6b that shows the non-numerical answers word’s frequency.

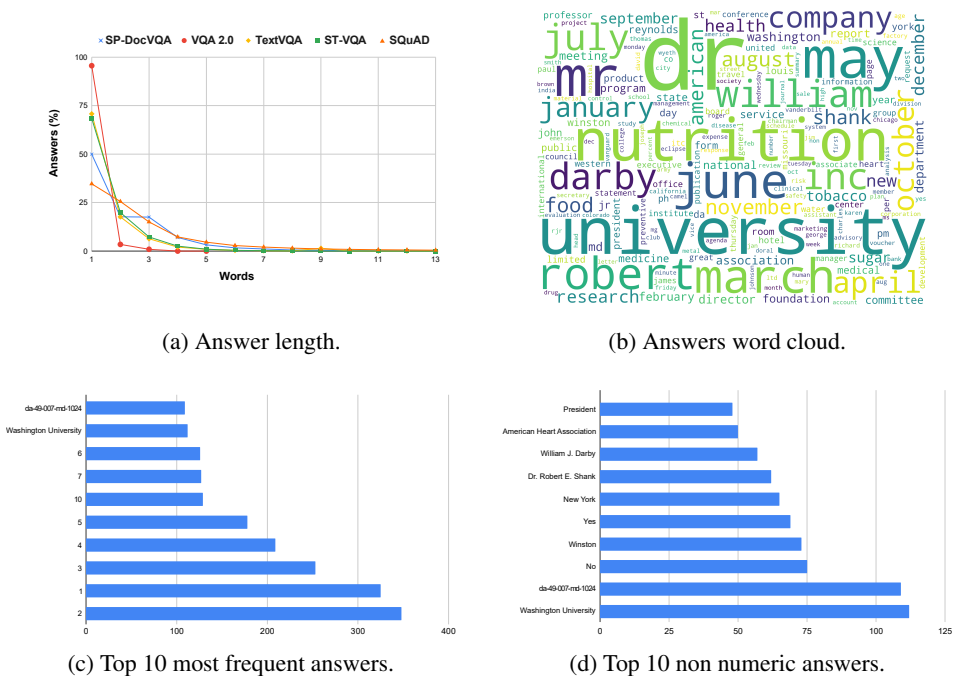


Figure 5.6: Answer statistics and comparison with similar VQA and MRC datasets. Most answers are numeric values and some named entities.

Figure 1 is a line graph showing the percentage of documents containing a given word across five different VQA datasets. The x-axis is labeled 'Words' and ranges from 0 to 100. The y-axis is labeled 'Documents / Contexts (%)' and ranges from 0 to 15. The legend indicates five datasets: SP-DocVQA (blue diamonds), VQA 2.0 (red circles), TextVQA (yellow diamonds), ST-VQA (green squares), and SQuAD (orange triangles). The graph shows that for all datasets, the percentage of documents containing a word drops sharply as the word index increases, with most words appearing in less than 1% of the documents. VQA 2.0 and ST-VQA show the highest initial frequencies, while SP-DocVQA shows the lowest.

[illegible]

Figure 5.7: Number of text instances in the datasets OCR and context. The word cloud showing the most frequent OCR words coincides with some of the questions (dates, amounts) and answers (nutrition, university)

5.3 Baselines

To assess the task complexity, we evaluate different heuristics, MRC and ST-VQA models on the introduced SP-DocVQA dataset.

5.3.1 Heuristics and Upper Bounds

The heuristics we evaluate are:

- **Random answer:** Performance when a random answer is picked from the available answers in the train split.
- **Random OCR token:** Performance when a random OCR token from the given document image is picked as the answer.
- **Longest OCR token and Biggest OCR Token:** Measures the performance when the longest OCR token in terms of amount of characters, and the OCR token with the bounding box with the largest area are selected as answers.

We also provide the following upper bounds:

- **Vocab UB:** Performance when correctly predicting the answers from a fixed dictionary comprising all the answers that occur more than once in the train split.
- **OCR subsequence UB:** Upper bound of predicting the correct answer from the subsequence of the image OCR tokens'. Therefore, the OCR tokens must be in the same exact order as the answer.
- **OCR substring UB:** Upper bound of predicting the correct answer from any combination of the recognized OCR tokens.
- **Human performance:** We collect the answers for all the questions in the test set with the help of few volunteers from our institution. Similar to the instructions given to the remote workers who annotated the dataset, we asked the volunteers to enter the answer extracted verbatim from the document text.

5.3.2 VQA Models

To evaluate the performance of existing VQA models on SP-DocVQA, we employed the two only models that, at the time of publishing the dataset, considered the text present in the images: LoRRA [191] and M4C [71].

LoRRA: follows a bottom-up and top-down attention [5] scheme, with additional bottom-up attention over the OCR tokens from the images. In LoRRA, the question tokens are first embedded using the pre-trained embedding GloVe [159], which are iteratively encoded using an LSTM [69] encoder. The model uses two types of spatial features to represent the visual information from the images (i) grid convolutional features from a Resnet-152 [65] pre-trained on ImageNet [40] and (ii) features extracted from bounding box proposals from a Faster R-CNN [171] object detection model pre-trained on Visual

Genome [104]. OCR tokens from the image are embedded using the pre-trained word embedding FastText [20]. An attention mechanism is used to compute a weighted average attention over the image features and the OCR tokens' embeddings. These averaged features are combined and fed into a classification output module. The classification layer of the model predicts an answer either from a fixed vocabulary (made from answers in the train set), or from a dynamic vocabulary, which is essentially the list of OCR tokens in the image. Here, the copy mechanism can copy only one of the OCR tokens from the image and consequently, it cannot output an answer which is a combination of two or more OCR tokens.

M4C: Was the current the state-of-the-art method on ST-VQA, TextVQA, and OCR-VQA datasets at the time of this work. It uses a multimodal transformer and an iterative answer prediction module. In M4C the tokens in questions are embedded using a BERT model [42]. The images are represented using (i) appearance features of the objects detected using a Faster-RCNN pretrained on Visual Genome [104] and (ii) location information as bounding box coordinates of the detected objects. The OCR tokens recognized from the image are represented using (i) the pretrained word embedding FastText, (ii) the appearance feature of the token's bounding box from the same Faster R-CNN used for the appearance features of objects, (iii) the PHOC [3] representation of the token and (iv) the bounding box coordinates of the token. Then these feature representations of the three modalities (question tokens, objects and OCR tokens) are projected into a common, learned embedding space and fed into a stack of Transformer [211] layers, where the multi-head self attention in the transformers enables both intra- and inter-modality attention. Finally, the answers are predicted through iterative decoding in an auto-regressive manner from a fixed vocabulary or from the recognized OCR tokens. Note that in this case the fixed vocabulary consists of the answer words, and not the answers itself as in the case of LoRRA.

In our experiments, we use the original LoRRA and M4C model implementations and few variant setups to adapt it to documents. Since document images in SP-DocVQA usually contain higher number of text tokens compared to images in scene text VQA datasets, we use larger dynamic vocabularies (i.e. more OCR tokens are considered from the images) for both LoRRA and M4C. Moreover, for both models, we also evaluate its performance when no fixed vocabulary is used. On the other hand, given that the notion of visual objects in real word images is not directly applicable in document images, we also try out variants where objects features are omitted.

5.3.3 Reading Comprehension Models

In addition to the VQA models which can read text, we evaluate extractive question answering models from NLP. In particular, we use **BERT** [42] question answering model, which is a pre-trained language representations from unlabeled text using Transformers [211]. These pretrained models can be used for downstream tasks with just an additional output layer. In the case of BERT, this is achieved by attaching a language layer at the end to predict the start and end indices of the answer span within the provided context.

5.4 Experiments

To evaluate the method’s performance, we use as the primary metric the Average Normalized Levenshtein Similarity (ANLS) introduced in ST-VQA, Chapter 3, section 3.2.4. We also provide the Accuracy for better interpretation of the results.

5.4.1 Experimental setup

In all the experiments including heuristics and trained baselines, the OCR tokens are extracted using the commercial OCR Microsoft Azure [146]. For the heuristics, upper bounds, and VQA models we use a fixed vocabulary of the 5,000 most frequent answers in the train split.

For LoRRA and M4C models we use the official implementations available as part of the MMF framework [188], and use the same settings and hyperparameters as the ones reported in the original works. For BERT, we use two pre-trained models of different sizes from the Huggingface Transformers library [223], which we refer to as BERT-Base and BERT-Large. Moreover, we also show the performance when BERT is trained on SQuAD before fine-tuning on SP-DocVQA. In extractive question answering or reading comprehension datasets, ‘contexts’ on which questions are asked are usually passages of electronic text. Instead, in SP-DocVQA ‘contexts’ from document images are simulated by concatenating the OCR tokens into a single string following the top-left to bottom-right reading order. In addition, to fine-tune the BERT QA models the answers are formatted as a ‘span’ of the contexts, defined by the start and end indices of the answer. To approximate the answer spans, we follow the approach proposed in TriviaQA [88], which consist of finding the first match of the answer string in the serialized OCR string.

5.4.2 Results

The results of all heuristic approaches and upper bounds are reported in Table 5.1, where it can be seen that all heuristics fail to achieve even 1% of accuracy in the validation or test splits. Notably, the performance when choosing the Longest or Biggest bounding boxes achieve worse performance than the comparable STR(bbox) in ST-VQA (Table 3.2), that achieves 0.128 of ANLS in the open dictionary task. This suggests that the intuition that humans tend to formulate questions about the largest text instance in the image is not as frequently applicable in the context of documents. Documents, being more textually dense and featuring particular layouts, might direct human attention over other regions.

Baseline	Val		Test	
	ANLS	Acc.	ANLS	Acc.
Random answer	0.003	0.00	0.003	0.00
Random OCR token	0.013	0.52	0.014	0.58
Longest OCR token	0.002	0.05	0.003	0.07
Biggest OCR token	0.010	0.62	0.011	0.67
Vocab UB	-	31.31	-	33.78
OCR subsequence UB	-	76.37	-	77.00
OCR substring UB	-	85.64	-	87.00
Human performance	-	-	0.981	94.36

Table 5.1: Evaluation of heuristics methods and upper bounds on SP-DocVQA. Predicting random answers or majority answer do not even yield 1% accuracy. Answers are a substring of the serialized OCR output in more than 85% of the cases.

The Vocabulary Upper Bound shows that the variability of the answers in the dataset allows to correctly answer only the 33.31% and 33.78% of the questions in the validation and test splits. This highlights the need of extra dynamic vocabularies or mechanisms that allows to include the image text as answers. The *OCR subsequence UB* shows that the answer is a sub-sequence of the concatenated OCR for around 76% of the questions in both validation and test splits, while *OCR substring UB* shows that the combination of the recognized words in any order are the correct answer for around 86% of the questions. This means that for around 10% of the questions, the words forming the answers are not consecutively positioned in the recognized OCR sequence, which will require the methods to reason about different places in the document. Moreover, since SP-DocVQA is proposed as an extractive task, and all words of the answers must be present in the image, it means that for 14% of the questions, the words comprising the answer have been recognized with errors, or they have not been recognized at all.

The results of the VQA baselines are reported in Table 5.2. The first row in both methods reports the results of the models following the original settings proposed by their respective authors. In case of LoRRA, the original setting proposed by the authors yields to the best results compared to the proposed variants we try out. With no fixed vocabulary, the performance of the model drops sharply, suggesting that the model primarily relies on the fixed vocabulary to output the answers. Larger dynamic vocabulary results in a slight performance drop, suggesting that incorporating more OCR tokens from the document images makes the attention over the dynamic vocabulary more sparse, and becomes more difficult to train the model. Unlike LoRRA, M4C benefits from a larger dynamic vocabulary, since increasing the size of the dynamic vocabulary from 50 to 500 improves by around 10 points of ANLS. In addition, when object features are omitted, M4C performs slightly better compared to the original setting, which is expected as recognized objects have little sense in documents.

Method	Objects' features	Fixed vocabulary	OCR Tokens	Val		Test	
				ANLS	Acc.	ANLS	Acc.
LoRRA [191]	✓	✓	50	0.110	7.22	0.112	7.63
	✓	✗	50	0.041	2.64	0.037	2.58
	✗	✓	50	0.102	6.73	0.100	6.43
	✓	✓	150	0.101	7.09	0.102	7.22
	✓	✓	500	0.094	6.41	0.095	6.31
M4C [71]	✓	✓	50	0.292	18.34	0.306	18.75
	✓	✗	50	0.216	12.44	0.219	12.15
	✗	✓	50	0.294	18.75	0.310	18.92
	✗	✓	150	0.352	22.66	0.360	22.35
	✗	✓	300	0.367	23.99	0.375	23.90
	✗	✓	500	0.385	24.73	0.391	24.81

Table 5.2: Performance ST-VQA models on SP-DocVQA. Detection of visual objects and their features (bottom-up attention), a common practice in VQA is ineffective in case of SP-DocVQA.

Results of the BERT question answering models are reported in Table 5.3, which shows that BERT, in any of the settings perform better than the best VQA baseline, which is M4C (last row in Table 5.2). This suggests that modelling the long textual context in documents is closer to machine reading comprehension tasks, as the number of entities involved in ST-VQA task are usually much lesser than in documents. As expected, the best performing BERT setting is the BERT-Large when, trained on SQuAD and SP-DocVQA, performing ~ 5 points of ANLS better than BERT-Large, which suggests, as in other similar works, that the amount of data used strongly boosts the results obtained.

Model	Fine-tune on SQuAD	Fine-tune on SP-DocVQA	Val		Test	
			ANLS	Acc.	ANLS	Acc.
BERT-Base	✗	✓	0.556	45.60	0.574	47.60
BERT-Large	✗	✓	0.594	49.28	0.610	51.08
BERT-Large	✓	✗	0.462	36.72	0.475	38.26
BERT-Large	✓	✓	0.655	54.48	0.665	55.77

Table 5.3: Performance of BERT question answering model. BERT-Large model trained on SQuAD [167] and fine-tuned on SP-DocVQA performs the best.

In Figure 5.8 we show the ANLS performance by question types, comparing the best VQA and BERT models against the human performance on the test split. While the human performance is uniform along different question types, BERT performs better in running-text and forms, where usually standard reading order with the textual semantics is enough to answer the questions. Instead, its performance drastically drops for figures and photographs, where spatial and visual features are very relevant. Moreover, BERT seems to perform bad on the handwritten type. This is interesting as the OCR recognition system is prepared for handwritten text and this drop might be caused by the position of the text, which are often handwritten annotations placed in ways that overlap with other words, are rotated, or appear between different lines. These factors can alter the standard reading order of the paragraph, and pose difficulties for accurate recognition. On the contrary, M4C poor performance doesn’t seem to stand out in any of the question types. In Figures 5.9 to 5.13 we show some qualitative results from our experiments, demonstrating a different behavior of BERT and M4C models on different question types.

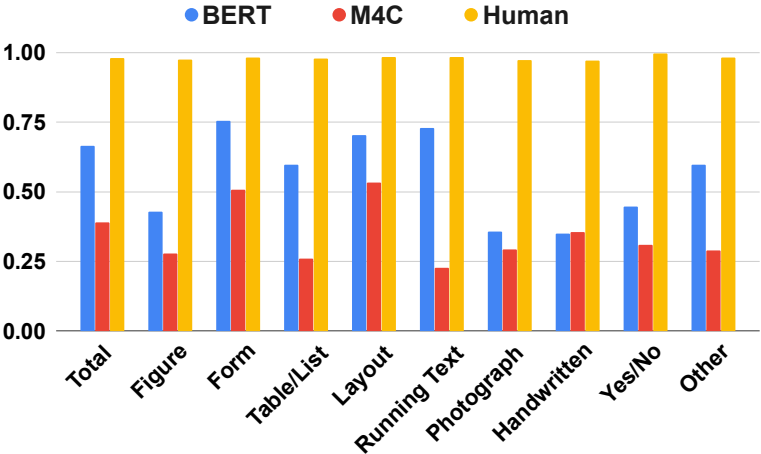


Figure 5.8: Best BERT and M4C baselines against the human performance for different question types. We need models which can better understand figures and text on photographs. Moreover, models that can better read and interpret handwritten text could improve the performance on this category.



Human: 9

Figure 5.9: Qualitative results from our experiments. The BERT model can’t answer correctly the leftmost example due to the lack of visual features, while M4C responds correctly. In the second example, the BERT model correctly answers a question on a form while the M4C model fails. In the rightmost example, both models fail to understand a step by step illustration.

island desserts (continued from page 7)

HAWAIIAN PUFF CAKE

- 1 cup butter, melted
- 1 cup packed brown sugar
- 1/2 cup dried cranberry orange peel
- 1/2 cup dried pineapple
- 1 1/2 cups chopped dates
- 1 cup packed cherries
- 1 cup dried coconut flakes
- 1 cup dried bread crumbs
- 1/2 cup dried mandarin slices
- 1 tablespoon butter
- 1 tablespoon sugar
- 1 teaspoon ginger juice
- 1/2 cup flour
- 1 cup shortening
- 1 cup brown sugar
- 1/2 cup raisins
- 1 teaspoon cinnamon
- 1 teaspoon salt
- 1/2 cup pineapples

Preheat oven to 375 degrees. Grease two 9" x 9" inch metal pans. Line each pan with brown paper. Cream butter and sugar in mixer. Sprinkle with honey, cherry and ginger juice. Stir in 1 cup of flour. Beat in shortening and brown sugar. Add remaining ingredients. Press into prepared pans. Bake for 20 minutes. Makes 2-2 1/2 dozen.

MALDAISANS

- 1 package yeast
- 1 cup warm water
- 1/2 cup brown sugar
- 4 cups flour
- 1/2 cup sugar
- 1 cup warm milk
- 1 1/2 cups melted butter
- 4 cups slightly brown

Preheat oven to water and 1 tablespoon sugar. Measure dry ingredients in a large bowl, and melted butter. Boil water, dissolved with 1/2 cup sugar, salted mixture. Stir in 1/2 cup brown sugar. Add 1/2 cup warm milk. Stir and stand until it starts to bubble. Form into balls. Bake for 10 minutes. Heat and drop into oil and cook until brown. Then drain on paper. Serve warm. Available in bulk.

HATPIA

- 2 cups coconut milk (filices)
- 1 cup sugar
- 4 tablespoons sugar
- 4 tablespoons cornstarch

Heat coconut milk in a double boiler. Combine and stir the ingredients separately until smooth. Cook and stir until thick. Add sugar and cornstarch. Stir until smooth. Add slightly and stir the pudding ingredients. Poured in from hot liquid. Cook until thick. Stand for 10 minutes. Stir and drop into oil. Let stand until firm. Then drain and serve.

FOR THE RECORD: The *Island* is published weekly on Thursdays, except on the 29th. Box 1044, Manassas, Virginia 20108

In a word...

EXOTIC

Based on a sailing ship off the coast...

Explores the exotic night skies...

Hawaii authentic, quaint and serene...

And discovers the life of ancient Hawaii.

Bishop Museum's

PASSPORT TO POLYNESIA

SEE YOUR HOTEL TRAVEL
AGENT OR VISIT
THEATRE KINGS TRAVEL
1-800-233-3388

Island from the DIAMOND PALACE
NEW LOCATION: World Travel, Park Mall
10000 W. International Blvd. Suite 1000 • Northridge • Ventura, CA 91324

48 *Letting the 1990s Begin*

island desserts (continued from page 7)

<p>HAWAIIAN FRUIT CAKE</p> <ul style="list-style-type: none"> 1 cup seedless raisins ½ cup seeded raisins 2/3 cup diced citron ½ cup diced candied orange peel ¼ cup diced pineapple 1 cup chopped dates 	<p>MALASADAS</p> <ul style="list-style-type: none"> 1 package yeast ½ cup warm water 1 tablespoon sugar 6 cups flour ½ cup sugar
---	--

Q: What is the first recipe on the page?

QT: Layout

GT: hawaiian fruit cake

M4C: island desserts (continued from cake

BERT: hawaiian fruit cake

Human: hawaiian fruit cake

Figure 5.10: Examples where BERT QA correctly answers questions other than ‘running text’ type: In the first question, the model gets the answer correct except for an extra space. In the second example, it answers correctly, which can be explained because BERT can interpret both the meaning of what a recipe is, and the order, even if it is 1-D.

16. Courses in which you wish to enroll: (Please read instructions carefully and check the appropriate boxes).

A.M. Schedule
 Fundamentals of Biostatistics M-T-W-Th-F
 If you have selected this course, do not select the following morning courses:
 Alternate
☒ Principles and Methods of Epidemiologic Research T-Th-S
 or
☐ Epidemiology of Occupational Hazards T-Th-S
 and
☐ Epidemiology of Cancer M-W-F
 or
☐ Epidemiology of Cardiovascular Diseases M-W-F

P.M. Schedule
 Fundamentals of Epidemiology M-T-Th-F
 If you have selected this course, do not select the following afternoon courses:
 Alternate
☒ Health Services Planning and Evaluation T-Th
 or
☐ Epidemiology of Injuries T-Th
 and
☐ Infectious Disease Epidemiology M-W-F
 or
☒ Epidemiology of Nutritional Diseases and Abnormalities M-W-F

17. Please make the following room reservation for me:
☒ Middlebrook Hall (Two in a room, twin beds) Arrive _____ Depart _____

18. Signature of Applicant *John H. Hargrett* Date *4/4/76*

19. I approve of this application *John H. Hargrett*
 Department Chairman or Advisor

Send this form with check for \$25.00, made payable to the University of Minnesota, to: Dr. Leonard M. Schuman, Program Director, Epidemiology Summer Session, Division of Epidemiology, Room 1-117 Unit A, Health Sciences Building, University of Minnesota, Minneapolis, Minnesota 55455.

APPLICATIONS MUST BE RECEIVED BY MAY 17.

Let Yourself Grow!

The University of Minnesota adheres to the principle that all persons shall have equal opportunity and access to facilities in any phase of University activity without regard to race, creed, color, sex, age or national origin. Under this principle, educational, cultural, social, housing, extra curricular and employment opportunities are available to all on an equal basis.



Q: What is written inside logo in the bottom of the document?

QT: Layout

GT: let yourself grow!

M4C: yourself grow!

BERT: <no prediction>

Human: let yourself grow!

Fulfilling the potential to win
 Inspired individuals share their story—what drove them to succeed against the odds and how they realized their goals.

"The launch of these new brands has been the joint effort of everyone, right across the company. From buying the land to the blend, the R&D, the processing, the packaging and finally the marketing, sales and distribution. All had to work in tandem, to achieve one goal. It's really true, success is about teamwork. It is about five fingers coming together to form a fist."

<GPI grew by 8 times the industry rate of growth in 2002-08
 Successful launch of Prince and Prince

PRINCE FILTER

"I am really proud to be part of the team that worked on re-launching Prince in the market. We knew that we had a winning brand on our hands. We were so clear about the aim, we knew that we had to find a way to achieve it. That's why I believe that some day Prince will..."

The passion to succeed: It's a value that runs right across the organization. Every individual is determined to realize their best potential to be a winner.

Source: <https://www.industrydocuments.ucsf.edu/docs/znw023>



Q: What Tobacco brand of GPI is shown in the picture?

QT: Photograph

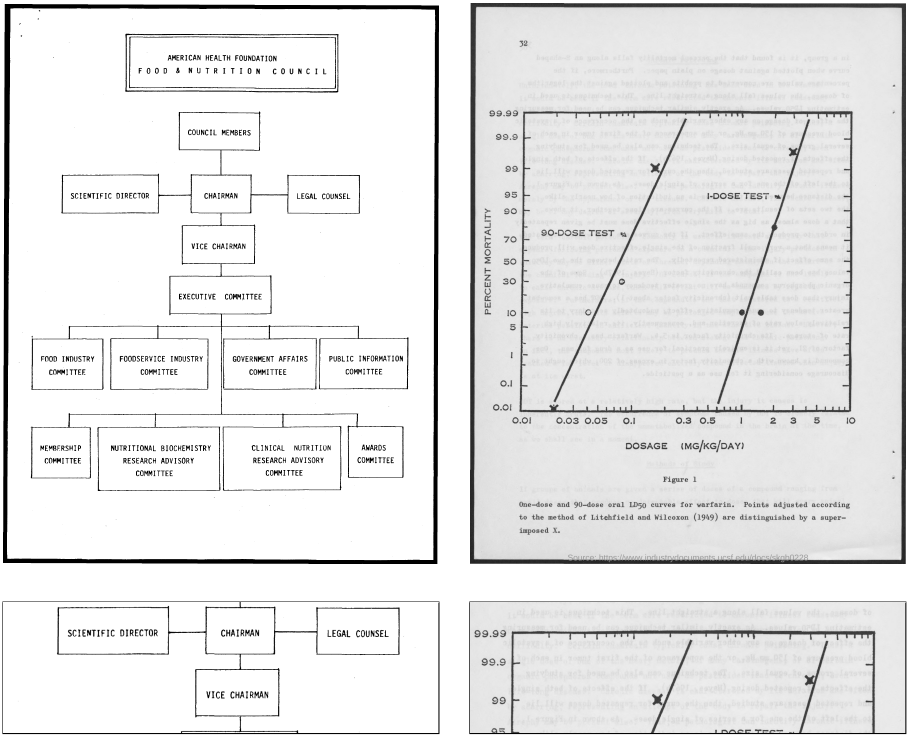
GT: prince

M4C: prince

BERT: <no prediction>

Human: prince

Figure 5.11: Qualitative results on questions based on pictures or photographs. BERT can't answer questions about text within logos or pictures, since the lack of visual features makes it unfeasible to contextualize such text. In contrast, M4C which extract the visual features of the OCR bounding boxes can correctly answer both examples.



Q: What is the position above "vice chairman"?

QT: Figure
GT: chairman
M4C: legal counsel
BERT: legal counsel
Human: chairman

Q: What is the highest value shown on the vertical axis?

QT: Figure
GT: 99.99
M4C: 50
BERT: 32
Human: 99.99

Figure 5.12: Understanding figures and diagrams. Both models fail to answer complex questions asked over diagrams. The structure and flow of an organizational hierarchy diagram or find the largest value of a ‘vertical axis’ seems to be too complex tasks for both methods.

Left Document (American Airlines Ticket):

Passenger: **larry** (GT: dr. william j. darby)
 Date: **1/7/77** (GT: 1/7/77)
 Flight: **740.00** (GT: 740.00)
 Class: **35.24** (GT: 35.24)

Right Document (Transmittal Memo):

To: **MICHELE** (GT: MICHELE)
 From: **CDF** (GT: CDF)
 Date: **1/7/77** (GT: 1/7/77)
 Subject: **File into "Nutrition" if we have one** (GT: File into "Nutrition" if we have one)
 Action: **3) Answer** (GT: Answer)

Q: What is the name of the passenger?
QT: Form, Handwritten
GT: dr. william j. darby
M4C: larry
BERT: larry
Human: dr. william j. darby

Q: What is the date present in the memo?
QT: Form
GT: 1/7/77
M4C: 1 7 7 7
BERT: 1 / 7
Human: 1/7/77

Figure 5.13: Impact of OCR errors. The accurate text recognition performance have a significant impact over the final performance. In these questions, both models are capable to correctly ground the questions over the relevant information, but failed to provide the correct answers only due to the OCR mistakes. In the first example, even the answer provided during the human performance study contains errors due to the complex handwriting style placed over the form line. In the second example, the OCR system has split the date into multiple tokens, resulting into incorrect answers of both the models.

5.5 Evolution of VQA on Single Page Documents

Since the publication of Single Page DocVQA there has been great progress in this topic, and we have received in the online evaluation platform more than 3190 submissions from 116 different teams, out of which 101 methods have been made public.




Figure 5.14: Evolution of method’s performance in SP-DocVQA dataset.

The winner of the DocVQA competition on single page document images was the DQA method [144] from the PingAn Chinese Insurance Company. It mainly consisted of a BERT-based transformer with 2-D positional embedding that was pretrained on Masked Language Modelling and answering synthetic generated questions. The representation of the spatial location in 2-D was very early adopted due to already proposed methods in ST-VQA and other document intelligence tasks. For example, LayoutLM [228] also used a 2-D positional embedding and pretrained the network to align both modalities. During the downstream task, the method also incorporated the visual features from the text bounding boxes to make use of this modality. Pre-training tasks were also widely adopted by most of the methods, each one adapting these strategies to align the used modalities. LayoutLMv2 [227] extended his predecessor by using different learnable bias to explicitly model the relative position, while they also make use of the visual features extracted from a ResNeXt [226] during the novel pretraining tasks proposed to align the three semantic, spatial and visual modalities. The last version, LayoutLMv3 [72], replaced the CNN image backbone for a simple linear embedding of the image patches, which turned the model much lighter and efficient. Another relevant contribution was brought by TILT [161], which extends the T5 [166] architecture, by replacing the positional embeddings with a 2-D spatial bias in the self-attention, and uses a U-NET to embed the visual features. Since T5 is a generative method, TILT became the first generative method for Document Visual Question Answering.

However, all these methods make use of an off-the-shelf OCR that is used beforehand, and the extracted text is fed to the proposed methods. In contrast, several methods such as SelfDoc [119], DocFormer [7], Donut [100], Dessurt [38] and Pix2Struct [112] have attempted to perform OCR-Free DocVQA, implicitly learning to read and interpret the document layout from the images only. However, although these methods perform comparable results compared to the rest of the methods on other Document Understanding tasks like Document Classification or Information Extraction, they show a significant performance gap in DocVQA, which suggests that they lack of a fine-grain semantic information necessary to answer the questions.

Since DocVQA and SP-DocVQA in particular is an interesting topic also for the industry, many companies have proposed their methods [7, 38, 72, 119, 161, 227]. Even though some of them lack of related paper such as Mybank-DocReader, Tencent Youtu, BAIDU-DI, Alibaba DAMO or ERNIE-Layout 2.0 due to both confidentiality and that they are often based on engineering rather than novel research, they coped the top leaderboard for long time. The last two models Alibaba DAMO and ERNIE-Layout 2.0 are specially interesting examples, since both consist of an ensemble of 30 models each, which even though might seem a basic approach, combining the output of 30 models might become complex. Unfortunately, since these methods don't have related papers, there isn't much more information about them.

On the other hand, following the trend to pretrain the model not only in self-supervised objectives, but also in multiple supervised tasks before fine-tuning on the downstream task, UDOP [199] achieved state-of-the-art results. Finally, LLM have also found their place in SP-DocVQA. PALI-X [28] and GPT-4 [156], are both top-scoring methods.

5.6 Conclusion

In this work, we have presented an important extension to the Visual Question Answering (VQA) domain, by performing the task over Documents, and thus starting the sub-field of Document Visual Question Answering (DocVQA). The proposed task aims to inspire a *purpose-driven* approach in document image analysis and recognition research by requiring models to answer questions related to the content of documents. To facilitate research in this area, we have introduced a new dataset for DocVQA that encompasses diverse challenges such as understanding complex layouts, form and table understanding, and recognition of different font styles.

Moreover, we have conducted a thorough analysis of the proposed dataset, using various baseline models, including OCR-based heuristics, state-of-the-art VQA models and NLP pre-trained language models. Our experiments demonstrate that the SP-DocVQA dataset is highly challenging and requires models to exploit the inherent structure of documents and the interactions between textual and visual modalities. We have shown that although existing models can provide competitive performance on some subsets of the dataset, we identify several cases where integrating more visual information, and understanding the structure of the document is needed besides just using the serialized recognized text.

Our results also highlight the need for novel approaches that can leverage the rich contextual information present in documents to improve performance. Furthermore, we have shown that this task requires a different formulation from previous ST-VQA, as documents contain much more text than natural images, and the visual features of document images are vastly different from natural images. This implies that models for SP-DocVQA must be able to handle the complex visual structure of documents and effectively combine the textual and visual modalities. Hence, future research should focus on developing models that can effectively reason over the structural and semantic information of documents and provide meaningful and diverse answers to document-based questions.

Furthermore, it is important to note that in the SP-DocVQA dataset, questions are explicitly designed to have their answer in the OCR tokens of the document. This type of questions are the ones expected to be tackled reasonably well by NLP methods upon a serialized or structured recognition output over the document. However, this type of task is relatively straightforward and more challenging questions are possible in the DocVQA domain.

After the presentation of this work that introduced VQA on documents, several works have utilized the dataset to enhance the performance on this task. The main focus has been to effectively model the semantic content, layout position and visual features of the documents, that has usually been tackled by proposing different pre-training tasks with increasing amount of data. Even though there is also the promising set of OCR-free methods that are capable to recover from inaccurately recognized text. The LLMs are again the ones who achieve the best performance.

After the introduction of this work, which pioneered VQA on documents, numerous other studies have utilized the dataset to enhance performance on this task. The primary challenge lies in effectively modeling the semantic content, layout position, and visual features of the documents. Researchers have addressed this challenge by proposing various pre-training tasks, often with increasing amounts of data. Additionally, there is a promising set of OCR-free methods capable of recovering from inaccurately recognized text. Notably, Large Language Models (LLMs) have consistently achieved the best performance in this context

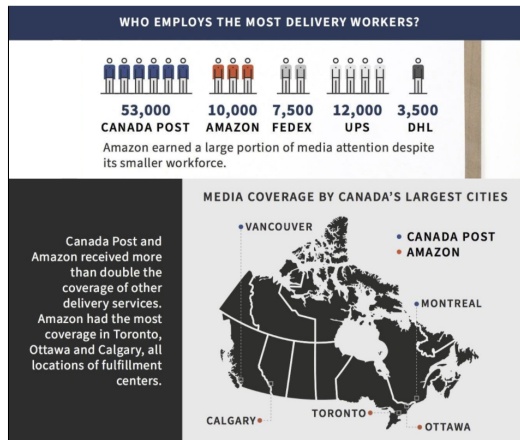
Chapter 6

Visual Question Answering on Infographic Images

Infographics communicate information using a combination of textual, graphical and visual elements. In this chapter, we explore the automatic understanding of infographic images by using Visual Question Answering. To this end, we present InfographicsVQA, a new dataset comprising a diverse collection of infographic images and question-answer annotations. The questions require methods to jointly reason over the document layout, textual content, graphical elements, and data visualizations. We create the dataset with an emphasis on questions that require elementary reasoning and basic arithmetic skills, and not all the questions are purely extractive. Furthermore, we evaluate two Transformer-based strong baselines that achieve unsatisfactory results compared to near perfect human performance on the dataset. The results suggest that VQA on infographic images, which are designed to communicate information quickly and clearly to the human brain, are ideal for benchmarking machine understanding of complex document images.

6.1 Introduction

Infographics are documents created to convey information in a compact manner using a combination of textual and visual cues. The presence of the text, numbers and symbols, along with the semantics that arise from their relative placements, make infographics understanding a challenging problem. Motivated by the multimodal nature of infographics and their human centered design, we propose a Visual Question Answering approach to infographics understanding as a mean to achieve better comprehension of documents with multimodal content. Thus, achieving comprehensive document image understanding in this domain is essential to answer questions such as the ones depicted in Figure 6.1, requiring the methods to jointly reason over the document layout, textual content, graphical elements, data visualizations, color schemes and visual art, among others.



Q: How many companies have more than 10K delivery workers?

Answer: 2

Answer evidence: Figure

Answer span: Non-extractive

Operation: Counting, Comparison

Q: Who has better coverage in Toronto - Canada post or Amazon?

Answer: canada post

Answer evidence: Text

Answer span: Question-span

Operation: None

Q: In which cities did Canada Post get maximum media coverage?

Answer: Vancouver, Montreal

Answer evidence: Text, Map

Answer span: Multi-span

Operation: None

Figure 6.1: Example image from InfographicsVQA along with questions and answers. For each question, span of the answer, type of evidence the answer is grounded on, and the discrete operations required to find the answer are shown.

Even though VQA has received significant attention over the past few years [5, 57, 60, 76, 96, 233], with several new VQA branches focusing on images with text, such as answering questions in business documents in chapter 5, text books [97], charts [26, 90, 91] and screenshots of web pages in VisualMRC [198]. Infographic images are unique in their combined use and purposeful arrangement of visual and textual elements. In Table 6.1 we show a summary of the latest VQA datasets and their main characteristics.

Dataset	Images	Synthetic Images	Template questions	Text type	# Images	# Questions	Answer type
TQA [97]	Science diagrams	✗	✗	MR	1K	26K	MCQ
RecipeQA [229]	Culinary pictures	✗	✓	MR	251K	37K	MCQ
ST-VQA [18]	Natural images	✗	✗	ST	23K	31K	Ex
TextVQA [191]	Natural images	✗	✗	ST	28K	45K	Ex, SAb
OCR-VQA [150]	Book covers	✗	✓	BD	207K	1M	Ex, Y/N
DVQA [90]	Bar charts	✓	✓	BD	300K	3.4M	Ex, Nm, Y/N
FigureQA [91]	Charts - 5 types	✓	✓	BD	120K	1.5M	Y/N
LEAF-QA [26]	Charts - 4 types	✓	✓	BD	250K	2M	Ex, Nm, Y/N
VisualMRC [198]	Webpage screenshots	✗	✗	BD	10K	30K	Ab
SP-DocVQA [143]	Industry documents	✗	✗	Pr, Tw, Hw, BD	12K	50K	Ex
InfographicsVQA	Infographics	✗	✗	BD	5.4K	30K	Ex, Nm

Table 6.1: Summary of VQA and Multimodal QA datasets where the text in the images is required to answer the questions. **Text type abbreviations:** Machine Readable (MR), Born Digital (BD), Scene Text (ST), Printed (Pr), Typewritten (Tw) and Handwritten (Hw). **Answer type abbreviations:** Multiple Choice Question (MCQ), Extractive (Ex), Short abstractive (SAb), Abstractive (Ab), Yes/No (Y/N), and Numerical - answer is numerical and not extracted from the image or question but derived from it (Nm).

In this chapter, we present InfographicsVQA, a new dataset for VQA on infographic images that includes 30,035 questions across 5,485 infographic images. The dataset contains questions grounded on tables, figures, and visualizations, along with questions that need multiple cues to answer. Moreover, as many infographics typically contain numerical data, the dataset also includes questions that require elementary reasoning skills such as counting, comparing numerical values, and arithmetic skills. However, in contrast to VisualMRC, the most similar dataset in terms of image and question types, InfographicsVQA is an extractive VQA task, except for questions that require certain discrete operations resulting in numerical non-extractive answers (section 6.2.2). With this dataset, it is possible to benchmark the advancement of algorithms at the intersection of vision, language, and document comprehension. We show an example of the InfographicsVQA dataset in Figure 6.1.

In addition, we leverage previous baselines used in SP-DocVQA BERT [42], and M4C [71], to show the performance difference between both datasets. And also fine-tune and adapt LayoutLM [228], a layout-aware document understanding method specifically designed for document understanding tasks, which we fine-tune for VQA on InfographicsVQA.

6.2 InfographicsVQA Dataset

6.2.1 Dataset creation

Infographic Images: Infographic images were crawled from the Internet by using the search query “infographics” on the Google and Bing image search engines. First, we removed duplicate images using a Perceptual Hashing approach implemented in imagedup library [82], and then, the annotation team described below filtered out undesired images by removing illegible or non-english images, non-infographic images or images without rich information and visual aspects.

Questions and Answers: The annotation process of questions and answers was carried out by a small number of trained annotators using an internal annotation tool. This strategy enabled us to control the quality of the questions, provide ongoing training for the annotation team, and shift the focus to different types of questions as needed. The annotation process consisted of two stages similar to the SP-DocVQA annotation process. In the first stage, workers were required to flag undesired images that were removed from the annotation platform, and then asked to write question-answer pairs based on the infographic images. In the second stage, they were instructed to answer the questions or flag them if they were unanswerable. Since the annotation was performed by the small annotation team, few question-answer pairs were considered wrong or ambiguous and hence, were discarded without the need of a third stage. At the end, the dataset was split into train, validation and test by 80% - 10% - 10% of the questions respectively.

6.2.2 Question-answer types

During the second annotation stage, in addition to answering the questions collected during the first stage, the annotators were instructed to add question-answer types (QA types) to better analyze and interpret the performance of different methods. In InfographicsVQA we collected the QA types from three different perspectives: Answer span, Answer evidence and Operations performed.

The **Answer span source** is defined based on where the answer can be found in the image, following the definition of ‘span’ in SQuAD [167] and SP-DocVQA. There are four types of answer-spans: Single-span, Multi-span, Question-span, and Non-extractive. A question-answer pair is considered as single-span if it corresponds to a single sequence of text tokens extracted verbatim from the image in the reading order. If the answer is composed of more than one span, is then tagged as multi-span. For example, in the last question of Figure 6.1, the answer is the name of two cities, which do not appear in a contiguous sequence of text. Therefore, it is considered as a multi-span answer. For multi-span answers, any order of the individual spans is considered a valid answer. In the previous example, both “Vancouver, Montreal” and “Montreal, Vancouver” are valid answers. Since such answers are unordered lists, we consider all permutations of the list as valid answers for the question at evaluation time. In addition, to increase the variabil-

ity of the questions, the annotators were also instructed to create multi-choice questions which are tagged as question-span. In this case, the answer to these questions is always provided within the question along with some distractors, like in the second question of Figure 6.1 or questions in Figures 6.10 and 6.13. Finally, the Non-extractive type is assigned when the answer is not text explicitly present in the image, and consequently, the answer must be generated. However, to limit the vast amount of language variability that can be considered correct for a given question, non-extractive questions were only allowed for numerical value answers such as the first question in Figure 6.1 and Figure 6.12.

The inclusion of multi-span, question-span, and numerical non-extractive answers in InfographicsVQA is a natural step in VQA involving text towards more complex visual question answering task different from the purely extractive setting in previous datasets ST-VQA and SP-DocVQA, while keeps a limited amount of possible correct answers in contrast to abstractive question answering in VisualMRC, where automated evaluation is difficult. As shown in section 6.4.2, near perfect human performance using automatic evaluation metrics (Table 6.3) validates that answers in InfographicsVQA are unique with limited variability when answered by different individuals.

The **Evidence type** indicates the kind of information in the document that is necessary to understand to answer the question. There are five types of evidences: Text, Figure, Table/List, Map and Visual/Layout. Map is used when the question is based on data shown on a geographical map, and Visual/Layout type is used when the evidence is based on the visual or layout aspect of the image. For example, questions such as “What is the color representing Canada Post in Figure 6.1, or “What is written in the top left corner?”.

Lastly, the **Operation type** captures the kind of discrete operation(s) required to infer the answer such as Counting, Arithmetic operations or Comparing numerical values. Although for most of the questions no further operation is required.

Figure 6.2 shows the distribution of questions based on their types. As evidenced in Figure 6.1, a question can have multiple answer sources, evidences, or operations, and many questions do not require any of the specified discrete operations to find the answer. For this reason, the total amount of questions in plots add up more than 100%.




Figure 6.2: Percentage of questions in validation and test set by their Answer span source (6.2a), Evidence required to answer (6.2b) and the discrete Operation necessary to find the answer (6.2c).

6.2.3 Summary, statistics, and analysis

The InfographicsVQA dataset comprises 30,035 questions posed over 5,485 images collected from 2,594 different web domains. In Table 6.2 we show basic statistics of questions, answers and OCR tokens in InfographicsVQA and other similar datasets and discuss it in the following section.

Dataset	Questions		Answers		Avg. Tokens per image	Avg. Image size (MP)	Avg. Aspect ratio
	%Unique	Avg. len	%Unique	Avg. len			
ST-VQA	84.84	8.80	65.63	1.56	7.52	0.35	1.43
TextVQA	80.36	8.12	51.74	1.51	12.17	0.77	1.41
VisualMRC	96.26	10.55	91.82	9.55	151.46	13.08	4.77
SP-DocVQA	72.34	9.49	64.29	2.43	182.75	3.91	1.35
InfographicsVQA	99.11	11.54	48.84	1.60	217.89	3.64	2.78

Table 6.2: Statistics of questions, answers and average tokens per image in InfographicsVQA and other scene text and document VQA datasets.

Questions: InfographicsVQA has the highest percentage of unique questions and the highest average question length compared to similar datasets. Moreover, the sunburst in Figure 6.3 displaying the first 4 words of the questions, shows that many questions inquire for quantities, specifically asking “How many” items or percentages. This is expected since infographic images often contain a significant amount of numerical data.




Figure 6.3: Starting 4-grams of questions in InfographicsVQA.

6.3 Baselines

In this section, we describe the baselines employed to evaluate the InfographicsVQA. These include heuristic baselines, upper bounds, and SoTA models for VQA and document image understanding.

6.3.1 Heuristic baselines and upper bounds

Heuristic baselines and upper bounds we evaluate are similar to the ones evaluated in previous SP-DocVQA.

Heuristic Baselines:

- **Random answer:** Performance when a random answer is picked from the available answers in the training set.
- **Random OCR token:** performance when a random OCR token from the given infographic image is selected as the answer.
- **Longest OCR token and Biggest OCR Token:** Measures the performance when the longest OCR token in terms of amount of characters, and the OCR token with the bounding box with the largest area are selected as answers.

We also provide the following upper bounds:

- **Vocab UB:** Performance upper bound on predicting the correct answer if the answer is present in the vocabulary composed of the most common answers in training set.
- **OCR subsequence UB:** Upper bound of predicting the correct answer from the subsequence of the image OCR tokens'. Therefore, the OCR tokens must be in the same exact order as the answer.
- **OCR substring UB:** Upper bound of predicting the correct answer from any combination of the recognized OCR tokens.
- **Vocab + OCR substring UB:** Combines the percentage of questions that are either in Vocab UB or OCR UB.
- **Human performance:** We collect the answers for all the questions in the test set with the help of few volunteers from our institution. Similar to the instructions given to the remote workers who annotated the dataset, we asked the volunteers to enter the answer extracted verbatim from the document text.

6.3.2 VQA Models

We employ previous VQA models **BERT** [42] and **M4C** [71] used in SP-DocVQA by fine-tuning and evaluate them on InfographicsVQA. This, not only allows to assess the performance of such methods on the new dataset, but also highlights the differences between the two datasets.

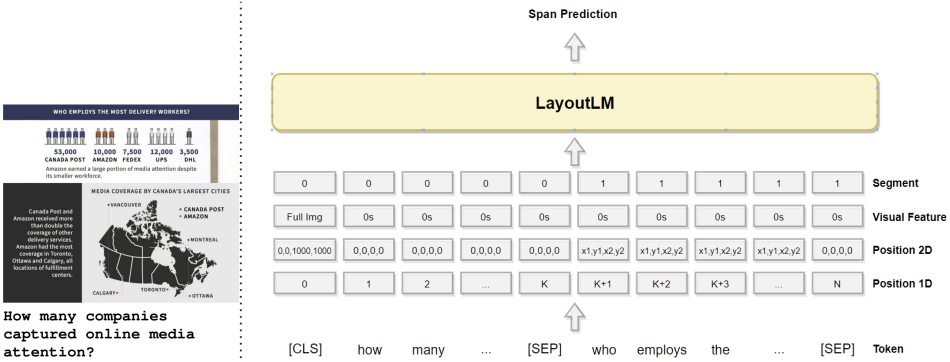


Figure 6.6: Overview of our LayoutLM based model for predicting answer spans. Textual, visual and layout modalities are embedded and mapped to the same space. Then, they are added and passed as input to a stack of Transformer layers.

Moreover, we also adapt **LayoutLM** [228], which proved to be effective for multiple document image understanding tasks [62, 73, 83] achieving state-of-the-art results at the time of this work, by extending BERT with layout information and visual features. In LayoutLM the input sequence starts with a special [CLS] token, followed by the question and OCR tokens. Between the question and OCR tokens a special [SEP] token is placed, which is further used at the end of the sequence to indicate its end. All the tokens in the input sequence are represented by its corresponding embedding which is the sum of i) a token embedding, ii) a segment embedding, iii) a 1D position embedding, iv) Four 2D position embeddings, and v) a visual embedding.

- **Token embeddings:** The token embeddings are the main core of LayoutLM, representing the semantic of the word. Following the original setting in BERT, the token embeddings, are represented with a 30,000 size vocabulary WordPiece [224] with hidden size of 768.
- **Segment embedding:** It is used to represent different segments or token sources in the input sequence. Consequently, the segment IDs representing the question and the OCR tokens are set to 0 and 1 respectively.
- **1D position embedding:** Is used to indicate the order of the tokens within the sequence. For the OCR tokens, we use the default reading order by concatenating them in top-left to bottom-right order.

- 2D position embedding: Are used to represent the position of the token within the image. Given the bounding box coordinates (x_1, y_1, x_2, y_2) of the OCR token, we embed all the 4 coordinate values using 4 separate embedding layers. x_1 and x_2 share same embedding table and y_1 and y_2 share another common embedding table. The coordinate values are normalized to lie in the range 0–1000 before embedding. The 2D embedding of the [CLS] token covers the whole image by setting it to (0, 0, 1000, 1000), while for the question tokens all the four 2D position embeddings are set to 0.
- Visual embedding: Unlike in original LayoutLM where visual features are fused after getting the attended embeddings from the Transformer block, we fuse the visual features early with the text input. Therefore, for each OCR token, we use the ROI pooled features from the Box Head of a pretrained object detection model. This feature vector is projected to 768 to match the size of the other embeddings and then, add up with the rest of the embeddings. For the [CLS] and other special tokens, we add the visual features corresponding to a ROI covering the entire image, named as “Full Img” in Figure 6.6.

Since LayoutLM is an extension of BERT, it is an extractive method which predicts the answer as a span within the provided context. Consequently, it can only answer questions whose answer span source is Single-span or Question-span, which represent nearly 70% of the questions. Extending this model to include questions with other Answer-sources like the fixed vocabulary of M4C, is an interesting direction for future work. A scheme of the LayoutLM-based model used for InfographicsVQA is shown in Figure 6.6.

Similar to the BERT and original LayoutLM, we train the model in two stages. First, we pre-train LayoutLM with the original Masked Visual-Language Model (MVLM) task using a masking probability of 0.15. Whenever masking, we replace each token with the [MASK] token 80% of the times, with a random token 10% of the times and keep it unchanged 10% of the times. Then, we fine-tune the model using the language head to predict the start and end token positions of the answer span.

6.4 Experiments

6.4.1 Experimental setup

In all the experiments, we use the text transcriptions and bounding boxes of the text tokens in the images extracted by the commercial OCR Amazon Textract [4]. For the heuristics, upper bounds, and VQA models we use a fixed vocabulary of the 5,000 most frequent answers in the train split. To evaluate the method’s performance, we use the ANLS introduced in section 3.2.4, and provide the Accuracy for better interpretation of the results.

In our experiments, for both M4C and LayoutLM models, we compare two different pre-trained object detection models for visual feature extraction: Faster-RCNN [171] trained on Visual Genome [104] (VG) and a Mask-RCNN [64] trained on document images in PubLayNet [244] for Document Layout Analysis (DLA). The Faster-RCNN model follows the original implementation provided in MMF framework [188] previously used in SP-DocVQA. The DLA detector is from the publicly available Detectron2 [225]-based implementation [68]. Features from the last Fully Connected (FC) layer of the Box head are used as visual features in M4C and LayoutLM models. In VG and DLA, these features are of size 2048 and 1024 respectively.

For the BERT and M4C methods, we use the official implementation of the BERT [143] and M4C models [188] previously employed in SP-DocVQA during Chapter 5. The fixed vocabulary of M4C is created from the 5,000 most common answer words in the train split. For LayoutLM, we start from the pretrained checkpoint made available by the authors in Huggingface’s Transformers [223], and the newly introduced linear projection layer, which maps the ROI pooled features to the common embedding size of 768, is initialized from scratch. Then, we pre-train the model on the Masked Visual-Language Model task on the train split of InfographicsVQA for 5 epochs, and finally fine-tune in VQA for another 10 epochs. To model the answers as spans within the context, we take the first subsequence match of the answer in the serialized transcription as the corresponding answer span.

6.4.2 Results

Results of heuristic baselines, upper bounds, and human performance are shown in Table 6.3. Following the same conclusions as in SP-DocVQA, none of the heuristic baselines is capable to correctly answer even 1% of the questions. On the other hand, the vocabulary upper bound allows to answer 20% more answers in InfographicsVQA than in SP-DocVQA (Table 5.1), because InfographicsVQA dataset comprises less unique answers as shown in Table 6.2. On the other hand, the upper bound provided by the OCR subsequence is 20% lower in InfographicsVQA, which is explained due to the amount of multi-span and non-extractive answers, and the inherent complex layout of infographic images that results in the serialization of the OCR tokens suboptimal for an extractive method. Moreover, the OCR substring UB shows a similar performance drop compared to SP-DocVQA due to the difficulty of OCR to recognize text printed with unique fonts and styles used in infographic images. When the Vocabulary and the OCR substring Upper Bounds are combined, it is possible to answer nearly 75% of the questions. Interestingly, even though there is a significant performance drop of the upper bounds and the VQA methods results described below between the SP-DocVQA and InfographicsVQA datasets, the human study on both datasets achieved similar results, suggesting that this task is not harder for humans.

Baseline	Val		Test	
	ANLS	Acc	ANLS	Acc
Random answer	0.006	0.00	0.005	0.00
Random OCR token	0.011	0.29	0.014	0.49
Longest OCR token	0.003	0.11	0.006	0.18
Biggest OCR token	0.009	0.25	0.009	0.21
Vocab UB	-	53.16	-	51.34
OCR subsequence UB	-	53.95	-	56.96
OCR substring UB	-	61.58	-	63.96
Vocab + OCR UB	-	76.71	-	77.40
Human performance	-	-	0.980	95.70

Table 6.3: Results of heuristics and upper bounds. Heuristics yield near zero results. More than 75% of the questions have their answer present either in a fixed vocabulary or as a Single-span of the OCR tokens serialized in default reading order.

The results using the M4C model in Table 6.4 shows an overall poor performance. Fine-tuning the visual features’ extractor does not affect the method performance, while the best strategy to integrate the visual features is to use a single ROI covering the whole image denoted as ‘Full Img’, regardless of using VG or DLA. Moreover, due to the high amount of text instances in the image, using only 50 OCR tokens from the image as in the original implementation is not enough, and the performance increases along with the size of the dynamic vocabulary. Finally, note that the M4C best performance is around 20 accuracy points less than in SP-DocVQA, demonstrating the increased difficulty of the task even for multimodal VQA methods.

Visual Features	Fine-tune detector	Object Features	OCR Tokens	Val		Test	
				ANLS	Acc	ANLS	Acc
VG	✓	100	50	0.107	4.81	0.119	4.87
VG	✓	20	50	0.111	4.82	0.122	4.87
VG	✗	20	50	0.125	4.89	0.127	4.89
VG	✗	20	300	0.128	4.90	0.134	5.08
VG	✗	None	300	0.136	5.86	0.143	6.58
VG	✗	(1) Full Img	300	0.142	5.93	0.147	6.64
DLA	✗	20	50	0.110	4.86	0.130	5.02
DLA	✗	20	300	0.132	5.95	0.144	6.50
DLA	✗	None	300	0.140	5.90	0.142	6.39
DLA	✗	(1) Full Img	300	0.138	5.97	0.140	6.42

Table 6.4: Performance of M4C. The original setting corresponds to the first row. ‘Fine-tune detector’ denotes if the last FC layer is fine-tuned along with the M4C model, which is the default setting in M4C. Object features column shows the amount of maximum objects (bottom-up features) used, using no object (‘None’) or using features from only one ROI corresponding to a box covering the entire image (‘Full Img’).

Results obtained with BERT reported in Table 6.5 show that strong language models can surpass VQA methods like M4C that considers layout position and visual features in complex images such as infographics. This is because the inherent world knowledge achieved by BERT during the pretraining stage can help to identify named entities, values, or dates. Moreover, it can guess some answers like the one in Figure 6.8 (center), where it is not necessary to look at the image to know that Oceania is the least populated continent. It is important to notice the performance drop of 0.460 in ANLS scored by BERT from 0.665 in SP-DocVQA (Table 5.3) to 0.205 in InfographicsVQA, which suggests that 2-D position and visual features are much more relevant in infographic images than in industry documents.

In-domain pretraining	Val		Test	
	ANLS	Acc.	ANLS	Acc.
✗	0.118	5.46	0.130	5.78
✓	0.203	12.14	0.205	12.26

Table 6.5: Performance of BERT and effect of extending the MLM pretraining task.

Results obtained with LayoutLM based model are shown in Table 6.6. In-domain pre-training helps the model significantly. Which is inline with the results of BERT and the observation by Singh *et al.* that pretraining on data similar to the data for a downstream task is highly beneficial in visio-linguistic tasks [189]. The best strategy to use the image features is to embed the ‘Full Img’ features from DLA, adding them to the CLS token, and using the visual features of each of the OCR bounding boxes. However, the use of visual features is in question due to the limited contribution to the overall performance. Indeed, the best accuracy in the test set is achieved by the LayoutLM setup using no visual features. This suggests that the iconographic domain of infographics is not well aligned neither with real scenes from Visual Genome nor classic documents from PublayNet.

Full Img to	Visual feature	In-domain pretraining	OCR visual	Val		Test	
				ANLS	Acc.	ANLS	Acc.
-	-	✗	✗	0.212	13.40	0.225	15.32
-	-	✓	✗	0.250	18.14	0.272	19.74
CLS	DLA	✓	✗	0.256	18.56	0.261	19.16
CLS	DLA	✓	✓	0.263	18.62	0.281	19.59
Non-OCR	DLA	✓	✓	0.245	17.21	0.263	18.37
CLS	VG	✓	✗	0.229	16.47	0.235	16.51
CLS	VG	✓	✓	0.118	16.71	0.115	16.94
Non-OCR	VG	✓	✓	0.042	1.75	0.037	1.28

Table 6.6: Performance of LayoutLM with different input settings. Row 1 and 2 show LayoutLM’s performance with and without in-domain pretraining. ‘Visual Feature’ column specify the kind of detector used for visual feature. ‘OCR visual’ indicate whether visual features of the OCR tokens are used or not. ‘CLS’ and ‘Non-OCR’ in ‘Full Img to’ column represent where the Full Image features are added: only to the CLS token or to all non OCR tokens (CLS, SEP and question tokens) respectively.

In Figure 6.7, we show the performance of the upper bounds, baselines and VQA methods per question types. The M4C and LayoutLM models used for this comparison are the variants that give the best ANLS on the test set. The most noticeable result is the low performance of all methods on the multi-span type of questions. M4C shows a consistent poor performance across all the categories, while LayoutLM performs better when the answer is purely extracted from single-spans of the text, and performs worse specially for questions that require to reason about visual components or maps, or requires performing any arithmetic operation. We also show qualitative results from the experiments in Figures 6.8 to 6.14.





Figure 6.7: Performance of baselines for different QA types.



Q: What is the interest rate of European Central Bank and US FED?


GT: 0.25%

M4C: 0.1%

BERT: 0.1%

LayoutLM: 0.25%

Human: 0.25%



Q: Which is the least populated continent in the world?

GT: Oceania

M4C: Oceania

BERT: Oceania

LayoutLM: EU

Human: Oceania



Q: What percentage of workers are not working from home?

GT: 23%


M4C: 66%

BERT: 77%

LayoutLM: 77%

Human: 23%

Figure 6.8: The left question evidence is Table/List, which only LayoutLM answers correctly. In the second question where evidence is a Table/List and it is required to compare different values, M4C answers correctly. BERT can also provide the correct answer, although it might be due to prior learned bias, since BERT can't capture the 2D position or color of the elements in the image. The last question requires to subtract 77 from 100, which none of the methods gets the answer correctly. For better visualization, images we show here are relevant regions cropped from the original infographics.



Q: What percent of researchers in Chile were men in the duration of 2011-15?

GT: [84%, 84]

M4C: 23%

BERT: 23%

LayoutLM: 23%



Human: 84%

Answer span: Single span

Answer evidence: Map

Operation: None


Figure 6.9: Color and geographic information requirements: To answer the depicted question, models need to understand that the green color corresponds to men, and then pick the number corresponding of the pie chart for Chile, but all baseline models answer the question wrong. Note that here there are two valid answers (GT) to account for language variability.



Q: Which quadrant does the country India fall into, blue, pink, or gray?
GT: pink
M4C: pink
BERT: economic activity
LayoutLM: country
Human: pink

Answer span: Question span
Answer evidence: Figure, Visual/Layout
Operation: None

Figure 6.10: The answer is a color name given as a multiple choice option in the question. To answer this question, the model needs to first locate “India” on the image and then identify the background color there. The multimodal aligned features of M4C allows it to correctly answer this question. In contrast, there is no text with ‘pink’ or any other color that extractive LayoutLM or BERT can use as answers and consequently, they output related concepts with India.



Q: Which are the top 2 best selling non violent video games of all time?

GT: super mario brothers, pokemon

Answer span: Multi-span

M4C: instagram youth

Answer evidence: Table/List

BERT: super mario brothers

Operation: Comparison


LayoutLM: super mario brothers

Human: pokemon, super mario brothers

Figure 6.11: The Multi-span answer type represent answers formed by multiple spans. In this example, the question requires finding the top 2 items within a category, which can only be answered if we allow answers containing multiple spans. Due to the extractive nature of the proposed LayoutLM-based model, it can only provide a part of the answer. Note that the permutations within the GT items scores perfect the answer provided in the human study.



Figure 6.12: Counting symbols/markers to find an answer: To answer this questions, the models need to count the number of yellow squares next to "CHAMPIONSHIPS". Since the answer is the result of a numerical operation, the non-extractive answer is allowed.



Q: Which drug was used more frequently during lockdown, MDMA, Cocaine, Cannabis, or Amphetamines?

GT: cannabis

M4C: cocaine

BERT: cannabis

LayoutLM: cannabis


Human: cannabis

Answer span: Question, Single

Answer evidence: Figure

Operation: Comparison

Figure 6.13: Comparing values shown in a bar chart: In this example, the answer is a span of the question (Question-span) and a span of the text on the image (Single-span) as well. To get the answer, it is necessary to compare the values of all the bars in the “Used more frequently” section. Hence, the operation to answer the question is tagged as ‘Comparison’.



Q: Playing against which country did he reach the most number of his milestone runs?
GT: sri lanka
M4C: pakistan
BERT: bangladesh
LayoutLM: bangladesh
Human: sri lanka

Answer span: Single span
Answer evidence: Figure
Operation: Arithmetic, Comparison

Figure 6.14: The context required to answer this question is spread along the entire image. The question requires the models to sum (Arithmetic) the number of milestone runs scored against each country, and then compare (Comparison) what is the opponent country to which the player has scored the most runs.

6.5 Evolution of VQA on Infographic Images

Since the release of the InfographicsVQA dataset, some researchers have used it to pre-train, train and/or fine-tune their methods on it. The dataset has also been used to evaluate the performance of some methods both in the standard or zero-shot settings. In total, we have received in the online evaluation platform more than 542 submissions from 31 different teams, out of which 30 methods have been made public. However, the specific field of infographic images has not received as much attention as the previous datasets. Most of the methods evaluated on the InfographicsVQA dataset are adapted methods employed for SP-DocVQA.




Figure 6.15: Evolution of method's performance in InfographicsVQA dataset.

The competition winner on InfographicsVQA was TILT, described earlier in section 5.5. Their main contribution on this dataset was to pretrain on a larger set of infographic images crawled on the Internet, by masking and reconstructing the words. The use of a U-Net [173] to extract visual features helped the model to better contextualize the words with the image icons, while the generative decoder allowed the model to predict answers out of the recognized words or vocabulary.

OCR-Free methods Donut [100] or Pix2Struct [112] have used the InfographicsVQA. However, they show similar performance gap as in SP-DocVQA, obtaining significantly worse results than methods using pre-extracted OCR. Finally, GPT-4 [156] has prevailed over the rest of the methods, achieving state-of-the-art results in zero-shot setting.

6.6 Conclusion

In this chapter, we have introduced the task of VQA on infographic images, creating the InfographicsVQA dataset and evaluating state-of-the-art methods on this novel task. While the result of the human study is similar to the achieved in SP-DocVQA, the results obtained using the strong baselines show that current state-of-the-art models for similar tasks perform poorly on the new dataset, which highlights the importance of correctly integrating the visual features to fully comprehend infographic images, and the need to create not only better VQA models that integrates all different modalities, but also devising improved feature extractors tailored specifically for processing infographic images. This has been latter approached by extending multimodal pretraining tasks by methods like TILT or GPT4, in which the domain-specific and vast amount of data used by each method respectively, helped both methods to better extract infographic image features, and integrate them with the rest of the cues.

Part III

Document collections and multipage documents

Chapter 7

Visual Question Answering on Document Collections

Tasks and methods in Document Understanding aims to process documents as single elements. However, documents are usually organized in collections (historical records, purchase invoices), that provide context useful for their interpretation. To address this problem, we introduce Document Collection Visual Question Answering (DocCVQA), a new dataset and related task where questions are posed over a whole collection of document images and the goal is not only to provide the answer to the given question, but also to retrieve the set of documents that contain the information required to infer the answer. Along with the dataset, we propose a new evaluation metric and baselines which provide further insights to the new dataset and task.

7.1 Introduction

Even though there has been a strong research effort on improving the machine understanding of documents, the research field of Document Analysis and Recognition (DAR) has focused on processing single page documents as in tasks like recognizing text from office documents [77], tables and forms [36], mathematical expressions [151] and visual information from figures and graphics [185]. Instead, the research on document collections has been limited to document retrieval by lexical content in word spotting [139, 168], blind to the semantics and ignoring the task of extracting higher level information from those collections.

On the other hand, over the past few years VQA has become one of the major relevant tasks as a link between vision and language. Moreover, although the ST-VQA dataset proposed in Chapter 3 start considering the text in the images in VQA, requiring the methods to reason about the recognized text to answer the questions, and the SP-DocVQA dataset proposed in Chapter 5 brought the task to industry documents, none of those works consider the image collection perspective, neither from real scenes nor documents.

In this chapter, we present Document Collection Visual Question Answering (DocCVQA) as a step towards better understanding document collections and going beyond word spotting. The objective of DocCVQA is to extract information from a document image collection by asking questions and expecting the methods to provide the answers. Moreover, to ensure that the answers have been inferred from the documents containing the necessary information and not from any prior learned bias, the methods are required to also provide the IDs of the documents used to obtain the answer as their evidence. Hence, we design this task as a retrieval-answering task, for which the methods should be trained initially on other datasets and consequently, we pose only a set of 20 questions over this document collection. In contrast to previous works, DocCVQA provides a purpose-driven approach and a collection wise point of view to Document Analysis and Understanding, posing new challenges and opening the door to bring new applications to industry. In addition, most of the answers in this task are actually a set of words extracted from different documents for which the order is not relevant, as we show in the question example in Figure 7.1. As a consequence, we define a new evaluation metric based on the Average Normalized Levenshtein Similarity (ANLS) to evaluate the answering performance of the methods in this task. Finally, we propose two baseline methods from different perspectives, which provide some insights on this task and dataset.

7.2 DocCVQA Dataset

In this section we describe the process to collect the images, questions and answers, show an analysis of the collected data, and describe the metrics used for the evaluation of the proposed task.

PUBLIC DISCLOSURE COMMISSION ode		Candidate Registration	C1 (1/2009)
Candidate's Name (Give candidate's full name.) ANNA M RIVERS			
Candidate's Committee Name (Do not abbreviate.) FRIENDS OF ANN RIVERS			
Mailing Address PO BOX 957			
City	County	Zip + 4	
LA CENTER	CLARK	98629	
1. What office are you running for?		Legislative District, County or City	Position No.
STATE SENATOR		LEG DISTRICT 17 - SENATE	N/A
2. Political party (if partisan office)		3. Date of general or special election	
REPUBLICAN		11-08-2016	
4. How much do you plan to spend during your entire election campaign, including the primary and general elections? Base the reporting options below. If no box is checked you are obligated to use Option II, Full Reporting. See instruction manuals for and changing reporting options.			
<input type="checkbox"/> Option I MINI REPORTING: In addition to my filing fee of \$ _____, I will raise and spend no more than \$5,000, including local voters pamphlets. I will not accept more than \$500 in the aggregate from any contributor except myself.			
<input checked="" type="checkbox"/> Option II FULL REPORTING: I will use the Full Reporting system. I will file the frequent, detailed campaign reports required by law.			
5. Treasurer's Name and Address. Does treasurer perform only ministerial functions? Yes ___ No <input checked="" type="checkbox"/> See WAC 390-05-243 and next page for details. List deputy treasurers on attached sheet.			
FRED RIVERS PO BOX 957, LA CENTER WA 98629			

PUBLIC DISCLOSURE COMMISSION ode		Candidate Registration	C1 (1/12)
Candidate's Name (Give candidate's full name.) ROBERT REEDY			
Candidate's Committee Name (Do not abbreviate.)			
Mailing Address PO BOX 61			
City	County	Zip + 4	
Mountlake TERRACE	SNO	98043	
1. What office are you running for?		Legislative District, County or City	Position No.
COUNTY CHARTER		SNO	4
2. Political party (if partisan office)		3. Date of general or special election	
NIP			
4. How much do you plan to spend during your entire election campaign, including the primary and general elections? Base the reporting options below. If no box is checked you are obligated to use Option II, Full Reporting. See instruction manuals for and changing reporting options.			
<input checked="" type="checkbox"/> Option I MINI REPORTING: In addition to my filing fee of \$ 0, I will raise and spend no more than \$5,000, including local voters pamphlets. I will not accept more than \$500 in the aggregate from any contributor except myself.			
<input type="checkbox"/> Option II FULL REPORTING: I will use the Full Reporting system. I will file the frequent, detailed campaign reports required by law.			
5. Treasurer's Name and Address. Does treasurer perform only ministerial functions? Yes ___ No ___ See WAC 390-05-243 and next page for details. List deputy treasurers on attached sheet.			
ROBERT REEDY			

Question: In which years did Anna M. Rivers run for the State senator office?

Answer: [2016, 2020]

Evidence: [454, 10901]

Figure 7.1: Partial visualization of sample documents in DocCVQA. The left document corresponds to the document with ID 454, which is one of the relevant documents to answer the question below. In DocCVQA the evidences are the documents where the information to answer the question is found. In this example, the correct answer are the years 2016 and 2020, and the evidences are the document images with IDs 454 and 10901 which corresponds to the forms where Anna M. Rivers presented as a candidate for the State senator office.

7.2.1 Dataset creation

Images: The DocCVQA dataset comprises 14,362 document images sourced from the Open Data portal of the Public Disclosure Commission (PDC), an agency that aims to provide public access to information about the funding of political campaigns, lobbyist expenditures, and the financial affairs of public officials and candidates. We choose this source for various reasons. It's a live repository that is updated periodically and therefore, the dataset can be increased in size in the future if necessary. In addition, performing VQA on this collection of documents can replicate a hypothetical realistic information extraction application among a set of scanned documents, by searching specific information of a candidate or party. Moreover, along with the documents, they provide their transcriptions in the form of CSV files, which allows to create questions and get their answers without the costly process of annotation. From the original collection of document images, we removed all the multipage documents and those which had incomplete

or ambiguous transcriptions. Thus, all documents that were finally included in the dataset consist of the same document template, the US Candidate Registration form, with slight design differences due to changes over the time. However, these documents still pose some challenges since the proposed methods will need to understand its complex layout, handwritten and typewritten text, while dealing with such large collection of documents. We provide some document examples in Figure 7.1.

Questions and Answers: To create the questions and answers, we identified the most relevant fields in the document form in terms of complexity (numbers, dates, candidate’s names, checkboxes and different form field layouts) and variability, listed in Table 7.1. We also defined different types of constraints (not, and, or...) since limiting the questions to find a specific value would place this in a standard word spotting scenario. Thus, we defined different conditions depending on the type of field required by the question: questions framed on date fields ask for documents before, after and between specific dates or years. For other textual fields, questions ask for documents with specific values (candidates from party P_1), for documents that do not contain specific values (candidates which do not represent the party P_1), or documents that contain a value from a set of possibilities (candidates from parties P_1 , P_2 or P_3). For checkboxes, we defined constraints regarding if a value is checked or not. Finally, according to the fields and constraints defined, the questions were posed in natural language referring to the whole collection, and asking for particular values instead of the document itself. Considering that DocCVQA is set up as a retrieval-answering task and documents are relatively similar among them, we pose only a set of 20 natural language questions over the collection which we manually inspect ensuring the correctness of the question-answer pairs, since some records of the document transcriptions provided are wrong. Finally, the questions were divided into two different splits; the sample set with 8 questions, and the test set with the remaining 12, which are shown in Table 7.2. Moreover, we ensured that in the test set there are questions asking for fields or containing constraints not present in the sample set. In addition, as depicted in Figure 7.2 the number of relevant documents is quite variable among the questions, which poses another challenge that methods will have to deal with.

Field	Type	# Values	# Unique values
Candidate name	Text	14362	9309
Party	Text	14161	10
Office	Text	14362	43
Candidate city	Text	14361	476
Candidate county	Text	14343	39
Election date	Date	14362	27
Reporting option	Checkbox	14357	2
Treasurer name	Text	14362	10197

Table 7.1: Documents form fields with an analysis of their variability, showing the number of absolute and unique values in the annotations.

ID	Question
8	Which candidates in 2008 were from the Republican party?
9	Which candidates ran for the State Representative office between 06/01/2012 and 12/31/2012?
10	In which legislative counties did Gary L. Schoessler run for County Commissioner?
11	For which candidates was Danielle Westbrook the treasurer?
12	Which candidates ran for election in North Bonneville who were from neither the Republican nor Democrat parties?
13	Did Valerie I. Quill select the full reporting option when she ran for the 11/03/2015 elections?
14	Which candidates from the Libertarian, Independent, or Green parties ran for election in Seattle?
15	Did Suzanne G. Skaar ever run for City Council member?
16	In which election year did Stanley J Rumbaugh run for Superior Court Judge?
17	In which years did Dean A. Takko run for the State Representative office?
18	Which candidates running after 06/15/2017 were from the Libertarian party?
19	Which reporting option did Douglas J. Fair select when he ran for district court judge in Edmonds? Mini or full?

Table 7.2: Questions in test set.




Figure 7.2: Number of relevant documents for each question in the sample set (blue) and the test set (red). Note that the y-axis is in logarithmic scale.

7.2.2 Evaluation metrics

The ultimate goal of this task is the extraction of information from a collection of documents. In consequence, the main evaluation metric assesses the method's answering capabilities. However, as a retrieval-answering style task, we also evaluate the retrieval of the correct documents, providing an intermediate metric to better interpret the final results.

Retrieval of evidences: Following standard retrieval tasks [125] we use the Mean Average Precision (MAP) to assess the correctness of the positive evidences retrieved by the methods. We consider positive evidences the documents in which a part of the answer to a given question can be found.

Answers: Following previous text based VQA tasks, we use the Average Normalized Levenshtein Similarity (ANLS) which captures the model's reasoning capability while smoothly penalizing OCR recognition errors. However, in DocCVQA the answers are a set of items for which the order is not relevant, in contrast to common VQA tasks where the answer is a single string. Thus, we adapt this metric to make it suitable for unordered list of answers. We name this adaptation as Average Normalized Levenshtein Similarity for Lists (ANLSL), formally described in Equation 7.1. Given a question Q , the ground truth list of answers $G = \{g_1, g_2 \dots g_M\}$ and a model's list of predicted answers $P = \{p_1, p_2 \dots p_N\}$, the ANLSL performs the Hungarian matching algorithm to obtain a k number of pairs $U = \{u_1, u_2 \dots u_K\}$ where K is the minimum between the ground truth and the predicted answer lists lengths. The Hungarian matching (Ψ) is performed according to the Normalized Levenshtein Similarity (NLS) between each ground truth element $g_j \in G$ and each prediction $p_i \in P$. Once the matching is performed, all the NLS scores of the $u_z \in U$ pairs are summed and divided by the maximum length of both ground truth and predicted answer lists. Therefore, if there are more or less ground truth answers than the ones predicted, the method is penalized.

$$U = \Psi(NLS(G, P))$$

$$ANLSL = \frac{1}{\max(M, N)} \sum_{z=1}^K NLS(u_z) \quad (7.1)$$

7.3 Baselines

This section describes the two baselines employed in the experiments. Both baselines tackle the task in two different stages. First, they rank the documents according to the confidence of containing the information to answer a given a question and then, get the answers from the documents with the highest confidence. The first baseline combines methods from the word spotting and NLP Question Answering fields to retrieve the relevant documents and answer the questions. We name this baseline as *Text spotting + QA*. In contrast, the second baseline is an ad-hoc method specially designed for this task and data, which consists of extracting the information from the documents and mapping it in the format of key-value relations. This baseline represents the collection similarly to how databases do, and consequently we name it *Database*.

7.3.1 Text spotting + QA

Evidence retrieval: We apply a text spotting approach to retrieve documents containing similar words to those in the questions, ranking the documents according to a confidence. For this, we first run a Part Of Speech (POS) tagger over the question to identify the most relevant words in it, keeping only nouns, numbers and dates, and ignore the rest of the words. Then, as described in Equation 7.2, given a question Q , for each relevant word in the question $w_i^q \in Q$ we get the minimum Normalized Levenshtein Distance (*NLD*) [114] between all recognized words w_j^{ocr} in the document and the question word. At the end, the confidence of a document containing the answer to the given question c_d is computed as the average of the distances of all question words.

$$c_d = \frac{1}{|Q|} \sum_{i=1}^{|Q|} \min_{j=1}^{|OCR|} \{NLD(w_i^q, w_j^{ocr})\} \quad (7.2)$$

Note that removing stopwords alone is insufficient. For instance, in the question depicted in Figure 7.1, the verb ‘run’ is not classified as a stopwords but cannot be found in the relevant document. As a result, other non-relevant documents might be ranked higher if they include such words.

Answering: Once the documents are ranked, we make use of BERT [42] question answering model to get the answers from the documents that have been ranked with a confidence c_d above a predefined threshold τ . BERT is a task-agnostic language representation based on transformers [211] that can be later used in other downstream tasks. In this case, we use extractive question answering BERT model which predicts the answer as a text span from a context, usually a passage or paragraph by predicting the start and end indices on that context. However, there is no such context in the DocCVQA documents that encompasses all the textual information. Therefore, we follow the approach of previous chapters in SP-DocVQA [143] and InfographicsVQA [142] to build this context by serializing the recognized OCR tokens on the document images to a single string separated

by white-spaces following the top-left to bottom-right order. Then, following the original implementation of [42] we introduce a start vector $S \in \mathbb{R}^H$ and end vector $E \in \mathbb{R}^H$. The probability of a word i being the start of the answer span is obtained as the dot product between the BERT word embedding hidden vector T_i and S followed by a softmax over all the words in the paragraph. The same formula is applied to compute if the word i is the end token by replacing the start vector S with the end vector E . Finally, the score of a candidate span from position i to position j is defined as $S \cdot T_i + E \cdot T_j$, and the maximum scoring span where $j \geq i$ is used as a prediction.

7.3.2 Database approach

This baseline aims to showcase the performance of an ad-hoc method using heuristics and commercial software to achieve the best possible performance. Since obtaining a human performance analysis is near impossible because it would require to manually check more than $14k$ documents for each question, this baseline can be seen as a performance to beat in a medium-long term.

The database baseline also tackles the task in the same retrieval and answering stages. However, in this case the ranking of the results is binary, indicating if a document is relevant or not. For this, we first employ a commercial information extraction tool over the document collection, obtaining the text and the key-value relationships between the field names and their values, including checkboxes. This is followed by a post-processing to normalize all the dates parsing them into the same format, and correct possible OCR recognition errors for fields with low variability (field names, parties and reporting options) by assigning the closest match given the edit distance between the recognized text and the set of classes. Finally, we map the key-value pairs into a database-like data structure. To answer the questions, these are manually parsed into structured query language and the resulting query is run over the stored records. If all the constraints are met, that document is considered relevant. Finally, the requested values in the question are extracted from the records of the relevant documents.

Note that this is a very rigid method that does not allow any modification in the data and therefore, it cannot generalize at all. Moreover, it requires a pre-process step to parse the query from Natural Language to a Structured Query Language (SQL) which is currently done manually, which makes the method not scalable also.

7.4 Results

7.4.1 Evidences

To initially assess the retrieval performance of the methods, we first compare the text spotting approach using two different commercial OCR systems and the database approach. The MAP of the evidence retrieval reported in Table 7.3 show that the performance of Amazon Textract [4] on text spotting is better than Google OCR [56], and is the only one capable of extracting the key-value relations for the database approach. In consequence, we use this as the standard OCR for the rest of the text spotting baselines.

Retrieval method	MAP
Text spotting (google)	71.62
Text spotting (textract)	72.84
Database	71.06

Table 7.3: Performance of different retrieval methods.

The overall retrieval performance of the database approach is similar to the text spotting approach. However, looking at the MAP obtained for each question depicted in Figure 7.3, we can see that it performs better for all the questions except for number 11 and 17. In both cases the problem is due to a failure in predicting the link between the key and value of the fields, and consequently it cannot retrieve documents with the required values, which is one of the main drawbacks of such rigid methods. Moreover, small OCR errors prevent the database method to correctly retrieve some of the required documents like in question 14, where the handwritten city value ‘North Bonneville’ has been recognized as ‘North Boweville’. On the other hand, the questions where the database approach have significantly better performance are those that require finding in the documents not only a particular value, but also understanding complex constraints such as the ones described in section 7.2.1, like finding a document between two dates (question 9), after a date (question 18), documents that do not contain a particular value (question 12), or where several values are considered as correct (question 14).




Figure 7.3: Evidence retrieval performance for each question in the test set.

7.4.2 Answers

For the BERT QA method, we use the pretrained and fine-tuned weights employed in SP-DocVQA from the Huggingface Transformers library [223], and further fine-tune on the DocCVQA sample set for 6 epochs to adapt the model to the new question concepts and context format. Since the sample set is small, during the 6 epochs the model only sees around 40 samples, which is sufficient to improve the answering performance without harming the previous knowledge.

Given the collection nature of DocCVQA, the answers to the questions usually consist on a list of texts spans located in different documents considered as relevant. Hence, we consider a document to be relevant if the confidence provided from the retrieval method is greater than a threshold τ . For the text spotting retrieval approach, we set the threshold to $\tau = 0.9$ after empirically finding to be the best value on the sample set. For the database approach, given that the provided confidence is binary 0 or 1, we consider as relevant all positive documents ($\tau = 1.0$).

In the experiments, we use the BERT QA baseline to answer the questions over the ranked documents from the text spotting and database retrieval methods. But we only use the database method to answer the ranked documents from the same retrieval approach. As reported in Table 7.4 the latter performs the best. The main reason for this is that the wrong retrieval of the documents yield the methods incapable to find the necessary information to answer the questions. In contrast, having the key-value relations allows the database method to directly output the value for the requested field as an answer, while BERT needs to learn to extract it from a context that has partially lost the spatial information. Moreover, the field value might be far from the field name, losing the semantic connection between the key-value pair. To showcase the answering performance upper bounds of the answering methods, we also provide their performance regardless of the retrieval system, where the documents are ranked according to the test ground truth.

Retrieval method	Answering method	MAP	ANLSL
Text spotting	BERT	72.84	0.4513
Database	BERT	71.06	0.5411
Database	Database	71.06	0.7068
GT	BERT	100.00	0.5818
GT	Database	100.00	0.8473

Table 7.4: Baselines results comparison.

As depicted in Figure 7.4, BERT does not perform well when the question asks for candidate’s names (questions 8, 9, 11, 14 and 18). However, it has a better performance when asking about dates (questions 16 and 17) or legislative counties (question 10). On the other hand, the database approach is able to provide the required answer as long as the key-value relationships have been correctly recognized. Interestingly, none of the methods are can correctly answer the question 13 regardless of a correct retrieval. This question asks if a candidate selected a specific checkbox value. The difference here is that the answer is *No*, in contrast to the sample question number 3. Then, BERT can’t answer because it lacks of a document collection point of view, it lacks of the interpretation of visual features to recognize the checkbox, and as an extractive QA method, it would require to have a *No* in the document along with some context that could help to identify that word as an answer. On the other hand, the database method fails because of its logical structure. If there is a relevant document for that question, it will find the field for which the query is asking for, or will answer ‘*Yes*’ if the question is a Yes/No type.




Figure 7.4: Baselines answering performance per question on the DocCVQA test set.

7.5 Evolution of VQA on Document Collections

Since the release of the DocCVQA dataset, there has been some progress on it. In total, we have received in the online evaluation platform more than 75 submissions from 10 different teams, out of which 9 methods have been made public. The challenging size limited variability of the dataset have turned this task hard to tackle. Resulting in little progress on it, where most of the methods used are small adaptations of SP-DocVQA methods.

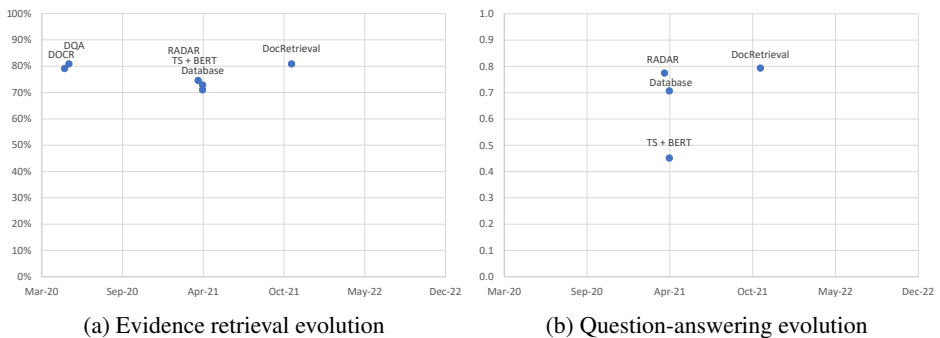


Figure 7.5: Evolution of method's performance in DocCVQA dataset

The challenge [144] held on this dataset was initially focused on the evidence retrieval by ranking the methods according to their MAP score, with the answering part as optional. Consequently, the methods PingAn DQA and iFLYTEK DOCR only focused on retrieving the relevant documents. The former used the same method employed for SP-DocVQA and keep the confidence scores to rank the methods, in contrast the latter proposed a custom fuzzy search algorithm based on the document text and layout features. The Infrd RADAR followed a similar approach as the proposed database baseline by extracting up to 28 fields from each document form, and parsed the language questions into SQL using the spaCy [70] library. Then, the method used the SQL queries with fuzzy-search algorithm to retrieve the set of relevant documents and the extracted answers. The current state-of-the-art method Mybank-DocRetrieval makes use of BERT method to train a classification task and use the logits as the ranking score to retrieve the documents. Then, they use the Mybank-DocReader proposed in SP-DocVQA to answer the questions.

7.6 Conclusions and Future Work

This work introduces a new and challenging task to VQA and DAR research fields. The DocCVQA that aims to provide a new perspective to Document understanding and highlight the importance and difficulty of contemplating a whole collection of documents. We have shown the performance of two different approaches. On one hand, a text spotting with an extractive QA baseline that, although it has lower performance, it is more generic and could achieve similar performance in other types of documents. And on the other hand, an ad-hoc baseline that represents the documents by their key-value relations that, despite achieving quite good performance, is still far from being perfect and because of its design is very limited and can't generalize at all when processing other types of documents. In this regard, we believe that the next steps are to propose a method that can reason about the whole collection in a single stage, being able to provide the answer and the positive evidences.

Chapter 8

Visual Question Answering on Multipage Documents

Existing work on DocVQA only considers single-page documents either alone or gathered in collections. However, in real applications documents are mostly composed of multiple pages that should be processed altogether. In this chapter, we propose a new multimodal hierarchical method Hi-VT5, that overcomes the limitations of current methods to process long multipage documents. In contrast to previous hierarchical methods that focus on different semantic granularity [63] or different subtasks [247] used in image classification, our method is a hierarchical transformer architecture where the encoder learns to summarize the most relevant information of every page and then, the decoder uses this summarized representation to generate the final answer, following a bottom-up approach. Moreover, due to the lack of multipage DocVQA datasets, we also introduce MP-DocVQA, an extension of SP-DocVQA where questions are posed over multipage documents instead of single pages. Through extensive experimentation, we demonstrate that Hi-VT5 is able, in a single stage, to answer the questions and provide the page that contains the answer, which can be used as a kind of explainability measure.

8.1 Introduction

Automatically managing document workflows is paramount in various sectors including Banking, Insurance, Public Administration, and the running of virtually every business. For instance, only in the UK more than 1 million home insurance claims are processed every year. Over the past 50 years, research on Document Image Analysis and Recognition (DIAR) has focused on specific information extraction and conversion tasks in order to automatize document workflows, which has placed the field at the meeting point between Computer Vision and Natural Language Processing. In the works presented in Chapters 5 to 7, we introduced the concept of Visual Question Answering (VQA) in DIAR, leading to different DocVQA tasks that require reasoning over typed or handwritten text, layout, graphical elements such as diagrams and figures, tabular structures, signatures and the semantics that these convey. This resulted in a paradigm shift, giving rise to end-to-end methods that condition the information extraction pipeline on a question posed in natural language.

However, all existing datasets and methods for DocVQA presented in previous chapters have focused on single page documents, which is far from real life scenarios. Documents are typically composed of multiple pages and therefore, in a real document management workflow all pages of a document need to be processed as a single set. In this chapter we aim to extend single-page DocVQA to the more realistic multipage setup.

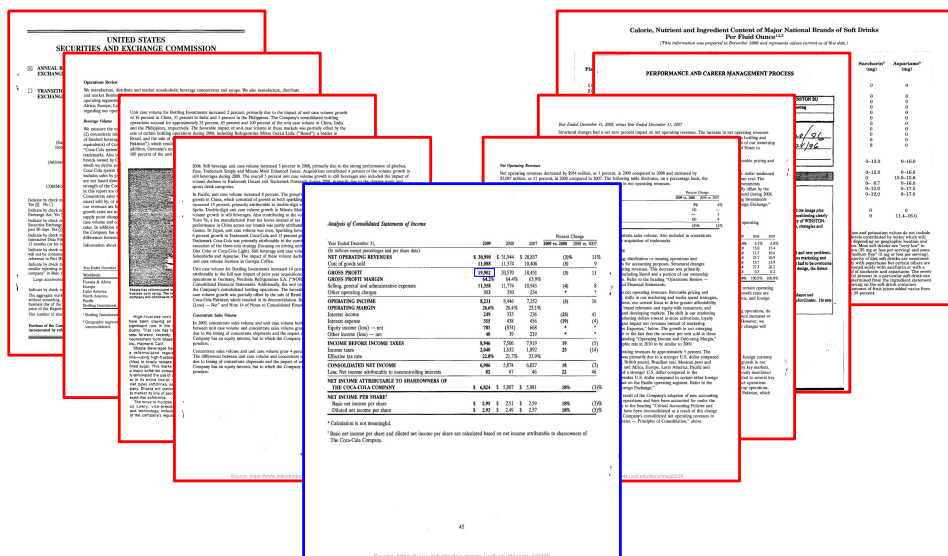


Figure 8.1: In the MP-DocVQA task, questions are posed over multipage documents where methods are required to understand the text, layout and visual elements of each page in the document to identify the correct page (blue in the figure) and answer the question.

One of the major challenges involved when processing multipage documents is the large increase of input data that models must process. DuReader_{vis} [162] is a Chinese dataset for open-domain document visual question answering that tackles a similar issue since each image contains 1968 tokens on average. The questions of this dataset are queries from the Baidu search engine, and the images are screenshots of the webpages retrieved by the search engine results. Although the answers are extractive, 43% of them are non-factual and much longer on average than the ones in previous DocVQA datasets. However, due to the big size of the image collection, the task is posed as a 2-stage retrieval and answering tasks similar to DocCVQA, where the methods must retrieve the correct page first, and answer the question in a second step.

The need to process longer sequences particularly affects current state-of-the-art DocVQA methods [72, 161, 227, 228] based on the Transformer architecture [211] that take as input textual, layout and visual features obtained from the words recognized by an OCR. As the complexity of the transformer scales up quadratically with the length of the input sequence, all these methods use a reduced number of input tokens which can lead to truncating a significant part of the input data when working with long multipage documents. To overcome this problem, several methods in the NLP domain have been proposed. Longformer [14] and Big Bird [236] replace the standard self-attention mechanism with a combination of global and local attention that reduces the computational requirements, which allows them to process longer sequences. However, they still present some limitations on the maximum sequence length they can accept, which degrades their performance on long multipage documents.

As an alternative, we propose the Hierarchical Visual T5 (Hi-VT5), a multimodal hierarchical encoder-decoder transformer built on top of T5 [166] which is capable to naturally process multiple pages by extending the input sequence length up to 20480 tokens without increasing the model complexity. In our architecture, the encoder processes separately each page of the document, providing a summary of the most relevant information conveyed by the page conditioned on the question. This information is encoded in a number of special [PAGE] tokens, inspired in the [CLS] token of the BERT model [42]. Subsequently, the decoder generates the final answer by taking as input the concatenation of all these summary [PAGE] tokens for all pages. Furthermore, the model includes an additional head to predict the index of the page where the answer has been found. This can be used to locate the context of the answer within long documents, but also as a measure of explainability, following recent works in the literature [204, 221]. Correct page prediction can be used as a way to distinguish which answers are the result of reasoning over the input data, and not dictated from model biases.

8.2 MP-DocVQA Dataset

The Multi-Page DocVQA (MP-DocVQA) dataset comprises 46K questions posed over 48K images of scanned pages that belong to 6K industry documents. The page images contain a rich amount of different layouts including forms, tables, lists, diagrams and pictures among others as well as text in handwritten, typewritten and printed fonts.

Dataset	Questions	Documents	Pages (Images)	Avg. pages per question	Document Avg. OCR Tokens
SP-DocVQA [143]	50K	6K	12K	1.00	151.46
VisualMRC [198]	30K	10K	10K	1.00	182.75
InfographicsVQA [142]	30K	5.4K	5.4K	1.00	217.89
DuReaderVis [162]	15K	158K	158K	1.3K	1968.21
DocCVQA [204]	20	14K	14K	14K	509.06
TAT-DQA [248]	16K	2.7K	3K	1.07	550.27
MP-DocVQA (ours)	46K	6K	60K	8.27	2026.59

Table 8.1: Comparison between MP-DocVQA and main DocVQA datasets.

8.2.1 Dataset creation

Documents naturally follow a hierarchical structure where content is structured into blocks (sections, paragraphs, diagrams, tables) that convey different pieces of information. The information necessary to respond to a question more often than not lies in one relevant block, and is not spread over the whole document. This intuition was confirmed during our annotation process in this multipage setting. The information required to answer the questions defined by the annotators was located in a specific place in the document. On the contrary, when we forced the annotators to use different pages as a source to answer the question, those become very unnatural and did not capture the essence of questions that we can find in the real world.

Consequently, we decided to use the SP-DocVQA dataset proposed in Chapter 5, which already has very realistic questions defined on single pages. To create the new MP-DocVQA dataset, we took every image-question pair from SP-DocVQA [143] and converted it to a multipage setting by adding the previous and posterior pages of the document downloaded from the original source UCSF-IDL¹, increasing the number of page images from 12,767 to 64,057, as shown in the second row of Table 8.2. As we show in Figure 8.2a most of the documents in the dataset have between 1 and 20 pages, followed by a long tail of documents with up to 793 pages. We focused on the most common scenario and limited the number of pages in the dataset to 20. For longer documents, we randomly selected a set of 20 pages that included the page where the answer is found.

¹<https://www.industrydocuments.ucsf.edu/>

Next, we filter out questions from the SP-DocVQA dataset that became ambiguous when posed in a multipage setup. We identified a set of key-words (such as 'document', 'image', 'graphic', etc.) that when used in a question on the multipage scenario could potentially lead to ambiguous answers. The most clear example is with the word 'document'. We can observe that in many documents most pages start with a big text on the top that can be considered as the title of the page. This would be the natural answer in the single page DocVQA scenario when the question asks about the title of the document. However, if this pattern is repeated in every page of the document, the question becomes impossible to answer when multiple pages are taken into account.

Moreover, even if there is only one page with a title, the answer can still be considered wrong, since the title of the document is always found in the first page like in the example in Figure 8.1. On the other hand, when we analyzed more closely other potentially ambiguous selected key-words such as 'image', 'appears' or 'graphic' we found out that the answers were not always ambiguous and also the amount of questions with those words was negligible compared to the entire dataset. Thus, we decided to keep those questions in our dataset. Finally, we found that the key-word 'title' was mostly ambiguous only when it was written along with the word 'document'. Hence, we decided to remove only the questions with the word 'document' in it, while keeping all the rest. This filtered version, which corresponds to the third row of Table 8.2 is the dataset version that has been released and used in the experiments, comprising 46,176 questions posed over 47,952 page images from 5,928 documents. Notice however, that the dataset also includes documents with a single page when this is the case, which corresponds to the fourth row. Nevertheless, the questions posed over multipage documents represents the 85.95% of the questions in the dataset.

	Documents	Pages	Questions
SP-DocVQA	6,071	12,767	50,000
MP-DocVQA (full)	6,071	64,057	50,000
MP-DocVQA (filtered)	5,928	60,884	46,176
MP-DocVQA (multi-page)	3,824	39,688	39,688


Table 8.2: Statistics of the MP-DocVQA during its construction process. In bold, we highlight the version (filtered) released and used for the experiments.

Finally, we split the dataset into train, validation and test sets keeping the same distribution as in SP-DocVQA. However, following this distribution some pages would appear in more than one split as they originate from the same document. To prevent this, we trim the number of pages used as context for such specific cases to ensure that any page is repeated between training and validation/test splits. In Figure 8.2b we show the number of questions according to the final document length.


To facilitate research and fair comparison between different methods on this dataset, along with the images and questions we also provide the OCR annotations extracted with Amazon Textract [4] for all the 64,057 document images.

8.2.2 Dataset statistics


As we show in Table 8.1, given that MP-DocVQA is an extension of SP-DocVQA, the average question and answer lengths are very similar to this dataset in contrast to the long answers that can be found in the open-domain datasets VisualMRC and DuReader_{vis}. On the contrary, the main difference lies in the number of OCR tokens per document, which is even superior to the Chinese DuReader_{vis}. In addition, MP-DocVQA adopts the multipage concept, which means that not all documents have the same number of pages (Figure 8.2b), but also that each page of the document may contain a different content distribution, with varied text density, different layout and visual elements that raise unique challenges. Moreover, as we show in Figures 8.2b and 8.2c the variability between documents is high, with documents comprising between 1 and 20 pages, and between 1 and 42,313 recognized OCR words.



(a) Document's pages before applying the 20 page limit



(b) Number of document's pages per question.



(c) Number of recognized OCR words per question

Figure 8.2: MP-DocVQA dataset statistics.

8.3 Baselines

As Multi-Page DocVQA is a new task, we adapt several state-of-the-art methods as baselines to analyze their limitations in the multipage setup and compare their performance against our proposed method.

BERT [42] is an extractive question answering method that creates the context for a given question by concatenating all the recognized words from an OCR. Then, the question and the context words are concatenated, tokenized and passed through the transformer layers and a classification head, which is used to predict the start and end indices of the answer within the context. Since BERT was the first question-answering method based on transformers, it shows the performance of such a simple baseline.

Longformer [14] is also an extractive question answering method based on RoBERTa [128] but it replaces the standard self-attention operation by a combination of global and local attention. The global attention is used on the question tokens, which attend and are attended by all the rest of the question and context tokens, while a sliding window guides the local attention over the context tokens to attend the other locally close context tokens. This attention operation scales linearly with the sequence length, reducing the algorithmic complexity from $O(n^2)$ to $O(n)$, and makes the model versatile for processing long documents and allowing sequence lengths up to 4,096 tokens.

Big Bird [236] extends the Longformer model by introducing random attention to the global and local attention. The random attention consists of setting global attention over randomly selected tokens, which will attend and be attended by all the rest of the tokens. This random attention improves the global understanding of the input sequence while only adding a marginal increase on the complexity of the attention pattern. In addition, Big Bird can work with two different strategies. The former, Internal Transformer Construction (ITC) only sets global attention over one single token, while the Extended Transformer Construction (ETC) sets the global attention over a set of tokens. Although the latter strategy is the desired setup for question-answering tasks by setting all the question tokens with global attention, the current released code only supports the ITC strategy and hence, we limit our experiments to this attention strategy.

LayoutLMv3 [72] is the current public state-of-the-art method on the SP-DocVQA dataset. It is the only baseline that uses visual features by representing the page image as patch embeddings. It also integrates a spatial-aware self-attention mechanism into the Transformer architecture to fully understand the relative 2D-positional relationship among different text blocks. The model is pretrained with 3 different objectives to align text, layout position and image context. Finally, a classification head is used for visual question answering to extract the answer from the given context.

T5 [166] is the first generative baseline and the backbone of the proposed method. The main difference of T5 is that it tackles all tasks with a text-to-text approach. Therefore, it does not output a label or a span of the input sentence, but the output is a text string. This, jointly with the new pre-training task of text denoising and huge amounts of data used to train it, has achieved very good results, reason why many current VQA state-of-the-art methods use it as backbone [15, 161].

However, all these methods are not directly applicable to a multipage scenario because they are designed to perform visual question answering on single-page documents or question answering from simple text streams. Moreover, they can't process all the document pages due to their limited input sequence length. Consequently, we define three different setups to allow them to be evaluated on this task. In the '*oracle*' setup, only the page that contains the answer is given as input to the transformer model. This setup aims to mimic the Single page DocVQA task, showcasing the raw answering capabilities of each model regardless of the input sequence length they can accept, by assuming the method has correctly identified the page containing the information. Thus, it should be seen as a theoretical maximum performance, assuming that the method has correctly identified the page where the information is found. In the '*concat*' setup, the context input to the transformer model is the concatenation of all the pages' contexts in the document. This can be considered the most realistic scenario where the whole document is given as a single input. It is expected that the large amount of input data becomes challenging for the baselines. The page corresponding to the predicted start index of the answer span is used as the predicted page, except for T5, since being a generative method it does not predict the start index. Finally, '*max conf.*' is the third setup, which is inspired in the strategy that the best performing methods in the DocCVQA challenge [204] use to tackle the big collection of documents. In this case, each page is processed separately by the model, providing an answer for every page along with a confidence score in the form of logits. Then, the answer with the highest confidence is selected as the final answer, with the corresponding page as the predicted answer page.

For BERT, Longformer, Big Bird and T5 baselines, we create the context following the standard practice of concatenating the OCR words in the image following the standard reading order: top-left to bottom-right. For all the methods, we use the Huggingface [223] implementation and pre-trained weights from the most similar task available.

8.4 Hi-VT5

Although documents contain dense information, not all of it is necessary to answer a given question. Following this idea, we propose the Hierarchical Visual T5 (Hi-VT5), a hierarchical encoder-decoder multimodal transformer where given a question, the encoder extracts the most relevant information from each page conditioned to the question and then, the decoder generates the answer from the summarized relevant information extracted from the encoder. Figure 8.3 shows an overview of the model. Each page is independently processed by the encoder taking as input the question, the recognized text and the image patches. Additionally, it incorporates a set of learnable $[PAGE]$ tokens, which are introduced to embed the relevant information of each page concerning a particular question at the output of the encoder. More specifically, the question and OCR words are first tokenized into individual subword tokens and embedded into the learned semantic contiguous representation, which captures the contextual information and semantics of the token. Moreover, the bounding box of the OCR tokens is represented by a continuous learned embedding space, which is latter added to the OCR semantic representation. In parallel, the page image is divided into non-overlapping patches, which are projected and fed to a Visual Transformer Encoder. Then a set of M $[PAGE]$ tokens are added into the input sequence length, which will be used to represent the summarized information conditioned by the question.

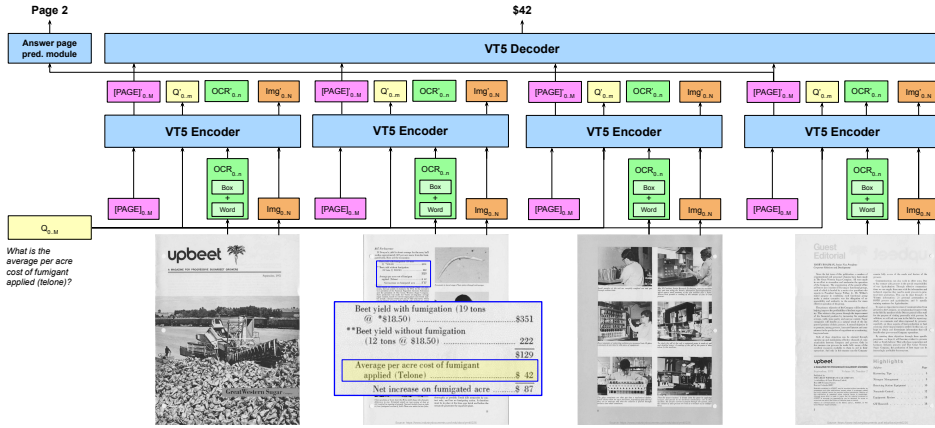


Figure 8.3: Architecture of Hi-VT5 method. Each page is processed by the encoder to represent in the contextualized $[PAGE]'$ tokens the most relevant information necessary to answer the posed question. Then, the $[PAGE]'$ tokens of all pages are concatenated to provide the decoder with a holistic representation of the document at the time of generating the answer. In addition, a classification layer in the Answer Page Prediction Module outputs the page where the answer to the question is found, providing the model with an explainability measure of the answers which allows, among others, to understand if the answer has been inferred from the actual input data, or from a prior learned bias.

All the tokens from the different sources: question, OCR, image patches and [PAGE] are concatenated and fed into the Hi-VT5 Encoder, which consists of a set of transformer layers, where self-attention mechanism is applied. Self-attention allows each token to attend to other tokens within the same or other modalities, capturing dependencies and relationships between the question, recognized words and image features. This allows the model to understand the content of each document page. The output of the Hi-VT5 encoder is a sequence of a high-dimensional representation for each token in the input sequence. Thanks to the pre-training and fine-tuning phase, Hi-VT5 learns to represent the necessary information to solve the task (either text de-noising or Visual Question Answering) in the contextualized [PAGE]' tokens at the output of the encoder. Following, the [PAGE] tokens are concatenated and passed through the decoder to generate the final answer in an auto-regressive way. Moreover, in parallel to the answer generation, the Answer Page Prediction Module (APPM) predicts the page index where the information to answer the question is found, which can be used as a kind of explainability measure.

We utilize the T5 architecture as the backbone for our method due to the vast amount of data and their novel de-noising task utilized during pretraining, which makes it an excellent candidate for the model initialization. Next, we describe each of the components of our pipeline.

Textual representation: Following recent literature on document understanding [72, 161] which demonstrates the importance of layout information when working with Transformers, we utilize a spatial embedding to better align the layout information with the semantic representation. Formally, given an OCR token O_i , we define the associated word bounding box as $(x_0^i, y_0^i, x_1^i, y_1^i)$. Following [15], to embed bounding box information, we use a lookup table for continuous encoding of one-hot vectors, and sum up all the spatial and semantic representations together:

$$\mathcal{E}_i = E_O(O_i) + E_x(x_0^i) + E_y(y_0^i) + E_x(x_1^i) + E_y(y_1^i) \quad (8.1)$$

where \mathcal{E}_i is the encoded representation for the OCR token O_i , and E_O , E_x and E_y are the learnable look-up tables.

Visual representation: We leverage the Document Image Transformer (DIT) [117] pre-trained on Document Intelligence tasks to represent the page image as a set of patch embeddings. Formally, given an image I with dimension $H \times W \times C$, it is reshaped into N 2D patches of size $P^2 \times C$, where (H, W) is the height and width, C is the number of channels, (P, P) is the resolution of each image patch, and $N = HW/P^2$ is the final number of patches. We map the flattened patches to D dimensional space, feed them to DiT, pass the output sequence to a trainable linear projection layer and then feed it to the transformer encoder. We denote the final visual output as $V = \{v_0, \dots, v_N\}$.

Hi-VT5 hierarchical paradigm: Inspired by the BERT [42] [CLS] token, which is used to represent the encoded sentence, we use a set of M learnable [PAGE] tokens to represent the page information required to answer the given question. Hence, we input the information from the different modalities along with the question and the learnable tokens to the encoder to represent in the [PAGE] tokens the most relevant information of the page conditioned by the question. More formally, for each page $p_j \in P = \{p_0, \dots, p_K\}$, let $V_j = \{v_0, \dots, v_N\}$ be the patch visual features, $Q = \{q_0, \dots, q_m\}$ the tokenized question, $O_j = \{o_1, \dots, o_n\}$ the page OCR tokens and $K_j = \{k_0, \dots, k_M\}$ the learnable [PAGE] tokens. Then, we embed the OCR tokens and question using Equation 8.1 to obtain the OCR \mathcal{E}_j^o and question \mathcal{E}^q encoded features. And concatenate all the inputs $[K_j; V_j; \mathcal{E}^q; \mathcal{E}_j^o]$ to feed to the transformer encoder. Finally, all the contextualized K' output tokens of all pages are concatenated to create a holistic representation of the document $D = [K'_0; \dots; K'_K]$, which is sent to the decoder that will generate the answer, and to the answer page prediction module.

Answer Page Prediction Module (APPM): Interpretability of the answer is an important research topic for the development of reliable VQA methods [221]. We use the prediction of the page where the answer is found as an evidence of the ability of the system to explain that the answer has been obtained using relevant information and not some prior bias. Thus, in parallel to the generation of the answer in the decoder, the contextualized [PAGE] tokens D are fed to a classification layer that outputs the index of the page where the answer is found.

Pre-training strategy: Since T5 was trained without layout information, inspired by [15] we propose a hierarchical layout-aware pretraining task to align the layout and semantic textual representations, while providing the [PAGE] tokens with the ability to attend to the other tokens. Similar to the standard de-noising task, the layout-aware de-noising task masks a span of tokens, replaces them with a single mask token, and forces the model to reconstruct the original set of tokens. Unlike the original de-noising task, in Hi-VT5 the encoder has access to the rough location of the masked tokens, which encourages the model to fully utilize the layout information when performing this task. In addition, the masked tokens must be generated from the contextualized K' [PAGE] tokens created by the encoder, which forces the model to embed the tokens with relevant information regarding the proposed task.

Training strategy: Even though Hi-VT5 keeps the same model complexity as the sum of their independent components: T5_{BASE} (223M) + DiT_{BASE} (85M) and despite being capable to accept input sequences of up to 20480 tokens, the amount of gradients computed at training time scales linearly with the number of pages, since each page is passed separately through the encoder and the gradients are stored in memory. Consequently, it is similar to have a batch size P times bigger in the encoder compared to a single page setting. While this could be tackled by parallelizing the gradients corresponding to a set of pages into different GPUs with different encoders sharing parameters, we offer a computational-efficient alternative strategy using limited resources. We train the model on shortened versions of the documents with only two pages: the page where the answer is found and the previous or posterior page. Even though this drops the overall perfor-

mance of the model, as we show in the ablation study (Table 8.7), training with only 2 pages is enough to learn the hierarchical representation of the model achieving results close to the ones using the whole document, and offers a good trade-off in terms of memory requirements. However, after the training phase, the decoder and the answer page prediction module can't deal with the full version of the documents of up to 20 pages. For this reason, we perform a final fine-tuning phase using the full-length documents and freezing the encoder weights.

8.5 Experiments

To evaluate the performance on this task, we use the standard evaluation metrics in DocVQA, accuracy and Average Normalized Levenshtein Similarity (ANLS). To assess the answer page prediction we use accuracy, i.e. if the predicted answer page is the correct one or not. For short, we name Answer Page Prediction Accuracy as APPA.

In Table 8.3 we show the baselines and Hi-VT5 hyperparameters employed to train the methods during the experiments. For the baselines' initialization weights, we choose the pretrained weights on the most similar task publicly available.

	BERT	Longformer	Big Bird	T5	Hi-VT5[†]
Model size	large	base	base	base	base
Parameters	334M	148M	131M	223M	316M
Model initial weights	SP-DocVQA	SQuADv1	TrivaQA	C4	C4
Max Seq. Length	512	4096	4096	512	20480
Training Loss	CE	CE	CE	CE	CE
batch size	32	8	8	32	8
lr	2e-5	1e-4	3e-5	2e-4	2e-4
optimizer	AdamW	AdamW	AdamW	AdamW	AdamW
scheduler	linear	linear	linear	linear	linear
warmup iterations	1000	1000	1000	1000	1000
training epochs	1	10	10	10	1 - 10 - 1

Table 8.3: Hyperparameters of the baselines and the proposed method that were used to train and evaluate on MP-DocVQA. †: Hi-VT5 refers to all three pre-training, training and fine-tune stages. The only difference is the number of epochs: 1, 10 and 1 respectively. Training loss CE denotes CrossEntropy loss.


8.5.1 Baseline results

As we show in Table 8.4, the method with the best answering performance in the oracle setup (i.e. when the answer page is provided) is T5, followed by LayoutLMv3, Big Bird, Longformer and BERT. This result is expected since this setup is equivalent to the single page document setting, where T5 has already demonstrated its superior results. In contrast, in the max conf. setup, when the logits of the model are used as a confidence score to rank the answers generated for each page, T5 performs the worst because the softmax layer used across the vocabulary turns the logits unusable as a confidence to rank the answers. Finally, in the concat setup, when the context of all pages are concatenated Longformer outperforms the rest, showing its capability to deal with long sequences. This can also be seen in section 8.5.1, which shows that the performance gap between Longformer and the rest of the baselines increases when the answer page is placed at the end of the document. The second best performing method in this setting is T5, which might seem surprising due to its reduced sequence length. However, looking at section 8.5.1 we can see that T5 is good on questions whose answers can fit into the input sequence, while it is not capable to correctly answer the rest. In contrast, Big Bird is capable to answer questions that require long sequences since its maximum input length is 4096, as Longformer. Nevertheless, it performs worse due to the ITC strategy, which does not set global attention to all question tokens and consequently, as long as the question and the answer tokens become more distant, it is more difficult to model the attention between the required information to answer the question.


Model	Size	Parameters	Max Seq. Length	Setup	Accuracy	ANLS	Ans. Page Accuracy
BERT [42]	Large	334M	512	Oracle	39.77	0.5904	100.00
				Max Conf.	34.78	0.5347	71.24
				Concat	27.41	0.4183	51.61
Longformer [14]	Base	148M	4096	Oracle	52.48	0.6177	100.00
				Max Conf.	45.87	0.5506	70.37
				Concat	43.91	0.5285	71.17
Big Bird [236]	Base	131M	4096	Oracle	55.39	0.6454	100.00
				Max Conf.	49.57	0.5854	72.27
				Concat	41.06	0.4928	51.78
LayoutLMv3 [72]	Base	125M	512	Oracle	58.81	0.6729	100.00
				Max Conf.	42.70	0.5513	74.02
				Concat	38.47	0.4538	51.94
T5 [166]	Base	223M	512	Oracle	59.00	0.6814	100.00
				Max Conf.	32.68	0.4028	46.05
				Concat	41.80	0.5047	—
Hi-VT5 (Ours)	Base	316M	20480	Oracle	50.01	0.6572	100.00
				Multipage	48.28	0.6201	79.23

Table 8.4: Baselines and proposed method Hi-VT5 results on MP-DocVQA dataset.

Baselines are evaluated on three different setups: oracle, concat and max conf. The proposed method is evaluated only on the oracle setup and the realistic multipage setting. We highlight in bold the best results for the oracle and any multipage (concat and max conf.) setup.



(a) Methods' answering performance by answer page position.



(b) Methods' APPA by answer page position.

Figure 8.4: Methods' performance by answer page position. In each subfigure we compare the performance of the baselines and Hi-VT5 in the 'oracle' setup (top). And the baselines on the 'max conf.' (middle) and 'concat' (bottom) setup against Hi-VT5 using its answer page prediction module. Note that the breakdown of the scores is not performed on the number of the document pages, but in which page the answer is located.

8.5.2 Hi-VT5 results

In our experiments we fixed the number of [PAGE] tokens to $M = 10$, through experimental validation explained in detail in the ablation study, section 8.6. We pretrain Hi-VT5 on hierarchical aware de-noising task on a subset of 200,000 pages of OCR-IDL [16] for one epoch. Then, we train on MP-DocVQA for 10 epochs with the 2-page shortened version of the documents and finally, perform the fine-tuning of the decoder and answer page prediction module with the full length version of the documents for 1 epoch. During training and fine-tuning all layers of the DiT visual encoder are frozen except the last fully connected projection layer.

Hi-VT5 outperforms all the other methods both on answering and page prediction in the concat and ‘*max conf.*’ setups, which are the most realistic scenarios. In addition, when looking closer at the ANLS per answer page position (see section 8.5.1), the performance gap becomes more significant when the answers are located at the end of the document, even compared with Longformer, which is specifically designed for long input sequences. In contrast, Hi-VT5 shows a slight performance drop in the ‘*oracle*’ setup compared to the original T5. This is because it must infer the answer from a compact summarized representation of the page, while T5 has access to the whole page representation. This shows that the page representation obtained by the encoder has still margin for improvement.

Identifying the page where the answer is found at the same time as answering the question allows to better interpret the method’s results. In Table 8.4 Hi-VT5 obtains the best APPA and, following the same behavior as in the answer generation, as section 8.5.1 shows the performance gap with the rest of the baselines becomes more significant as the answer page is located at the end pages of the document. Moreover, in Figure 8.5 we show the correlation between the performance of question-answering and answer page prediction, where it can be seen that our method can predict the correct page even when it cannot provide the correct answer. On the other hand, for some correctly answered questions, the corresponding answer page is wrongly predicted. This evidences that the answer has been inferred from a prior learned bias instead of the actual input data as shown in Figure 8.6 qualitative example.


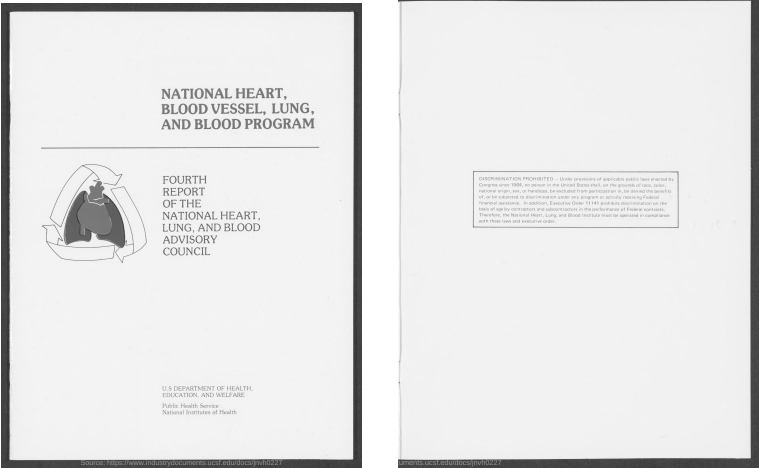


Figure 8.5: Matrix showing the Hi-VT5 correct and wrong answered questions depending on the answer page prediction module result.

To better understand the information that Hi-VT5 embeds into the [PAGE] tokens, we explore and visualize the attention scores for some examples in MP-DocVQA. The attention of Figure 8.7a, corresponds to the first [PAGE] token, which performs a global attention over the whole document with a slight emphasis on the question tokens, providing a holistic representation of the page. Other tokens like as shown in Figure 8.7c focus their attention over the other [PAGE], and question tokens. More importantly, there is always a token that focuses its attention to the provided answer, as seen in Figures 8.7b and 8.7d.



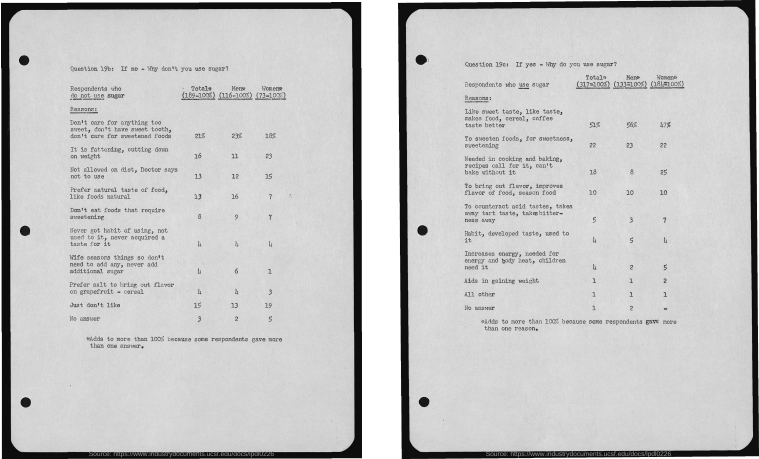
Question: In the image, what is the organ in the center?

GT Answer: **Heart**

Pred. Answer: **Heart**

GT Answer page: **0** (left)

Pred. Answer page: **13** (right)



Question: Based on this data, people of which gender (men or women) are more likely to avoid sugar to lose weight?

GT Answer: **Women**

Pred. Answer: **Women**

GT Answer page: **3** (left)

Pred. Answer page: **6** (right)

Figure 8.6: Using both metrics at the same time (ANLS and APPA) helps to interpret and filter answers that might have been inferred from prior learned bias instead of the actual input document. This Figure shows qualitative examples with correctly answered questions but wrong answer page predictions. (Top) The most common organ depicted in documents, and placed in the center of the human body makes the most plausible answer to be heart without looking at the documents. (Bottom) Binary question that can be answered with Men/Woman has a 50% chance of being randomly answered correctly. With the page wrongly predicted, it raises concerns regarding the source of the information used to answer the question.

PAGE] [PAGE] [PAGE] [PAGE] [PAGE] [PAGE] [PAGE] [PAGE] [PAGE]
- : : : : : : : : : :
approval form was filed ? context : : : : : when the

RJRT/RESEARCH & DEVELOPMENT
DOCUMENTATION APPROVAL FORM

Log-In No. 88-

NAME: William M. Coleman, III TYPE OF DOCUMENTATION: X
TITLE: Ex. RAD Chemist Abstract
DIVISION: Analytical Research Research Manuscript
DATE: 7/29/88 Other Presentation
MAIL ADDRESS: 611-132/002 Extension: 5177 Other RAD RADM or RADR Number & Date

DOCUMENT TITLE: "Minotaur Particulate Phase Comparison of a Reference Cigarette and a Cigarette That Heats Rather Than Burns Tobacco."

AUTHOR(S) (JCO-AUTHOR(S) : W. M. Coleman, III, W. L. Chinn, D. S. Moore, E. L. White, B. N. Gordon, M. S. Uhrig, J. A. Giles, J. F. Elder, Jr., M. F. Borden, and R. D. Hicks.

PURPOSE OF DOCUMENT

PUBLICATION IN:

PRESENTATION AT:

42nd Tobacco Chemists' Research Conference

LOCATION: Lexington, Kentucky Date: October 2-5, 1988

Other pertinent information: Abstract due no later than June 24, 1988.

Your signatures on this form indicate that: to the best of your knowledge, the information contained in the document to be published/presented is not proprietary and the technical quality reflects positively on the image of RJRT and its RAD Department.

Author(s): W. M. Coleman, III Date: 7/29/88
Manager: J. F. Elder, Jr. Date: 7-29-88
Director:

6688 RECIS

- | PAGE | PAGE | PAGE | PAGE | PAGE | PAGE | PAGE | PAGE |
|----------|----------|--------|------|---------|-------|-------|------|
| question | question | What | is | the | total | costs | |
| proposed | project | period | ? | context | | | |

SECTION 2 - **FINANCIALS** COMMUNICATION

HARVARD UNIVERSITY, ROOSE ALEX F. 270 56 3985

BUDGET ESTIMATES FOR ALL YEARS OF SUPPORT REQUESTED FROM PUBLIC HEALTH SERVICE								
DESCRIPTION	FISCAL YEAR 2000		FISCAL YEAR 2001		FISCAL YEAR 2002		FISCAL YEAR 2003	
	2000	2001	2002	2003	2004	2005	2006	2007
PERSONNEL COSTS	17,730	16,971	20,299	21,720				
CONSULTANT COSTS (CONSULTANTS, HON. FEE)	----	----	----	----				
EQUIPMENT								
SUPPLIES	400	400	400	400				
TRAVEL								
DOMESTIC	----	----	----	----				
FOREIGN	----	----	----	----				
PATENT COSTS								
ALTERNATING REVENUES	----	----	----	----				
Other Expenses	1,870	1,870	1,870	1,870				
Total Direct Costs	20,000	20,241	22,569	23,990				
Indirect Costs	6,300	6,764	7,313	7,718				
Total Costs	26,300	27,005	29,882	31,708				
TOTAL BUDGET REQUESTED PROJECT PERIOD (Show on Page 3, Item 6)				115,370				

REMARKS: Details of costs for 2000 and year for which 100% need rates are contained. For 2000 and 2001, support expenses include all work of support personnel for one full year. Following general method (continued) and is supported and approved. (See comments)

I, Valentin, Will select and plan content of the three investigations. Will direct **Will** writing and review **Will** literature.

A, Rose: will plan and supervise **Will** the statistical analysis of **Will** three topics.

C, Michael: will give and statistical analysis, coding and programming (statistician) under **Will** direction of Dr. Robert Reed.

Supplies: pens, pencils, paper, scanning.

Computer and Computer Time: coding and programming of data collected. **Starting** of program for projects.

END (FINALLY) PG 280

Page 252, January, 1990

Rev. 1/95

- Figure 8.7:** Visualization of the Hi-VT5 attention scores. Red indicates higher attention scores, yellow represents smaller attention scores, and transparent indicates no attention given at all

8.6 Ablation study

To validate the effectiveness of each of the components of Hi-VT5, we perform an ablation study and show the results in Table 8.5. Without the answer page prediction module the model performs slightly worse on the answering task, showing that both tasks are complementary and the correct page prediction helps to answer the question. The most significant boost comes from the hierarchical de-noising pre-training task, since it allows the [PAGE] tokens to learn better how to represent the content of the document. The last fine-tuning phase where the decoder and the answer page prediction module are adapted to the 20 pages maximum length of the MP-DocVQA documents, is specially important for the answer page prediction module. This is because the classification layer predicts only page indexes seen during training. Hence, without finetuning it can only predict the first or the second page. Finally, when removing the visual features the final scores are slightly worse, as in other works in the literature [15, 72, 161] which also show that the most relevant information is conveyed within the text and its position, while explicit visual features are not specially useful for grayscale documents.

Method	Accuracy	ANLS	Ans. Page Acc.
Hi-VT5	48.28	0.6201	79.23
–2D-pos	46.12	0.5891	78.21
–Vis. Feat.	46.82	0.5999	78.22
–APPM	47.78	0.6130	00.00
–Pretrain	42.10	0.5864	81.47
–Fine-tune	42.86	0.5263	55.74

Table 8.5: Hi-VT5 ablation study. We study the effect of removing different components independently from Hi-VT5 namely the 2D position embedding (2D-pos), visual features (Vis. Feat.), the answer page prediction module (APPM), the pretraining (Pretrain) and the last fine-tuning (Fine-tune) phase of the decoder and answer page prediction module.

We also explore the optimal number of [PAGE] tokens to embed the most relevant information from each page conditioned by the questions. Contrary to BERT [42], which represents a sentence with a single [CLS] token, Hi-VT5 requires more than one token to represent a whole page, since it conveys more information. The maximum number of tokens M that Hi-VT5 can process is limited by the decoder input sequence length S , and the number of pages P that must be processed. Formally,

$$M = \text{int} \left(\frac{S}{P} \right) \quad (8.2)$$

Therefore, we can set M as an hyperparameter to select depending on the number of pages we need to process, where in the extreme cases we can represent a single page with 1024 [PAGE] tokens, or a 1024 page document with a single token for each page.

Regarding the proposed setting on MP-DocVQA, where we limit the maximum number of pages per document to 20, the maximum possible number of tokens M would be 51. In consequence, we performed a set of experiments with different number of [PAGE] tokens to find the optimal value. As we show in Table 8.6, the model is able to answer correctly some questions even when using only one or two tokens. However, the performance increases significantly when more tokens are used. Nevertheless, the model does not benefit from using more than 10 tokens, since it performs similarly either with 10 or 25 tokens. Moreover, the performance decreases when using more. This can be explained because the information extracted from each page can be fully represented by 10 tokens, while using more, does not only provide any benefit, but also makes the training process harder.

[PAGE] Tokens	Accuracy	ANLS	APPA
1	36.41	0.4876	79.87
2	37.94	0.5282	79.88
5	39.31	0.5622	80.77
10	42.10	0.5864	81.47
25	42.16	0.5896	81.35
50	30.63	0.5768	59.18

Table 8.6: Results of Hi-VT5 with different [PAGE] tokens.

Finally, we explore the effect of training with different amount of pages to find a trade-off between the memory required and the performance achieved. In the results reported in Table 8.7, it can be seen that training with only one page, the performance is much lower compared to the others, suggesting that the model can't interpret the hierarchical paradigm. However, when training with more than two pages, although the performance keeps increasing, the boost is every time more limited, while the required memory keeps increasing.

Trained pages	Acc	ANLS
1	22.96	0.3860
2	33.37	0.5577
5	34.08	0.5730
10	34.25	0.5792

Table 8.7: Experiments showing the results when training with different number of document pages and tested with the document original length.

8.7 Conclusions

In this chapter, we have presented the task of Visual Question Answering on multipage documents and the related dataset MP-DocVQA. We have proposed the method Hi-VT5, that can process sequences up to 20,480 tokens and answer the questions regardless of the page in which the answer is placed without increasing the backbone model complexity, surpassing the evaluated baseline’s performance which includes methods specifically designed to process long sequences. In addition, the APPM can predict the page where the answer is located, which serves as an evidence of the answer and as an interpretability measure.

Chapter 9

Conclusions and Future Directions

9.1 Conclusions

Visual Question Answering is still a challenging task at the intersection of Computer Vision and Natural Language Processing. In this thesis, our goal has been to move towards more intelligent and holistic methods, integrating the semantics of the recognized text with the image visual features and question semantics in both natural scenes and document images, also contemplating multiple images at the same time. With this, we aimed at fostering research in interesting domains from both academic and industry perspectives, focusing on unexplored areas that will contribute to a broader comprehension of the real world and enable more effective utilization of AI methods in practical applications.

In chapter 3, we have shown the importance of the recognized text in the scenes by creating Scene Text VQA dataset, which has been widely used in the community after opening an entirely new research area. The dataset emphasizes the need to use the high-level features included in the images in the form of scene text to guide the VQA process, and includes highly variable questions and answers, which presents a challenge for existing VQA techniques that predicts the answer from the set of the most common answers. In addition, the proposed metric Average Normalized Levenshtein Similarity (ANLS), has proved to provide a smooth response to text recognition performance and to be more suitable for generative models than classification accuracy. In order to establish lower performance constraints and gain valuable insights, we have performed a series of tests with baseline approaches, showing the poor performance of illiterate methods. We have also shown that adding textual information to generic VQA models can improve their performance, but ad-hoc baselines (such OCR-based, which do include contextual terms) may outperform them, highlighting the necessity for better methods. The proposed method MMGrid enables to jointly reason about the textual and visual modalities in the scene by using an attention mechanism that pays attention to multi-modal grid characteristics. From the experimentation and ablation analysis, we have shown that focusing on multi-

modal characteristics is preferable to focus separately on each modality. The grid pattern we have chosen and the decision to use a one-stage item detection backbone rather than a classification one, turn out to be excellent choices for this task. Furthermore, we have demonstrated that the proposed model is adaptable enough to be paired with a traditional VQA model to produce cutting-edge results on mixed datasets that contain questions that cannot be directly addressed using OCR tokens.

In chapter 5, we have brought VQA to the documents domain, using natural language questions as an interface to unify different information extraction tasks, and providing an *end-purpose* to the document intelligence field driving the information extraction task on the required information. This has also opened another entire research line which have gathered a lot of attention from both the academic and industry sides. For this, we have created the SP-DocVQA dataset, and conducted a thorough analysis of the proposed dataset, showing that ST-VQA methods underperform on this dataset, which has images with many more text instances, and the fine-grained position and layout understanding is more relevant. Moreover, the bottom-up approach of the visual features may not be the best choice, as there is no element such as *objects* in documents, and this part of the methods becomes irrelevant. Instead, pretrained language models that focus on modelling the language and the interactions between the different words perform significantly better, even though they use simple 1D spatial position embedding and ignore the visual features. However, the impact of visual features on industry scanned documents is limited, and the dataset lacks of non-extractive questions, which can be seen as a limitation of the dataset.

In chapter 6, we have utilized infographic images to create the InfographicsVQA dataset, where the position and visual information are essential to answer the questions. This underscores the value of the dataset as an excellent tool for effectively evaluating the capabilities of multimodal methods in DocVQA. Through the experiments conducted, we demonstrated the limitations of state-of-the-art methods at the time, particularly in providing answers that required reasoning about multiple locations within the image or performing arithmetic operations. These type of questions and the intensive requirement of multimodal features made InfographicsVQA a more challenging dataset than SP-DocVQA. While we were able to demonstrate the value of spatial features in this task, we couldn't showcase the impact of visual features due to the complexity of extracting meaningful information from icons and colored entities. However, subsequent work, such as TILT [161], has validated the importance of visual information and modalities fusion during pre-training tasks, as employed in our adapted LayoutLM method.

In chapter 7, inspired by real applications where multiple documents are analyzed at the same time, we have applied VQA over a whole collection of single page documents, and created the retrieval-answering DocCVQA task and related dataset. Through the conducted experiments of the proposed methods, we have shown that simple word spotting along the different documents to find the relevant ones is not enough, since the complex questions posed in natural language contains complex clauses such as negation, conjunction or disjunction conditions. Moreover, we have shown that DocVQA methods still have a large gap for improvement, as an ad-hoc method based on pre-defined keys can outperform them.

In chapter 8, we have introduced the realistic multipage setting to DIAR research. For this, we have created the MP-DocVQA dataset to perform VQA on multipage documents in a single stage (end-to-end), by extending and curating the previous SP-DocVQA. The experiments conducted using thoroughly selected state-of-the-art methods revealed their limitations when processing large amounts of data spanning through multiple pages within a document, observing that even methods especially designed to handle long sequences, such as Longformer or Big Bird, struggle to answer questions located on the latter pages of the documents. In contrast, our proposed Hi-VT5, which can process up to 20,480 tokens without increasing the model’s complexity, outperforms the rest of the methods with a consistent performance regardless of where the answer is. This proves the excellent choice of setting the encoder and decoder of the model in a hierarchical paradigm, and the use of the [PAGE] tokens to embed the most relevant information of each page given a question. Moreover, the answer page prediction module of the Hi-VT5 can help to identify the source of the errors in incorrectly answered questions or possible learned biases.

Overall, we have demonstrated the crucial role of text in understanding natural scenes, and we have employed the VQA paradigm to showcase its potential as a natural language interface for document intelligence tasks. Furthermore, we have highlighted the significance of 2D spatial representation in these tasks, which becomes increasingly relevant as the layouts become more complex, allowing for a better association among closely located or particularly aligned words. Additionally, we have explored the application of VQA on multiple images, highlighting the necessity of creating methods capable to efficiently process large document collections, which would result into significant time savings for users utilizing its application in real-world scenarios.

9.2 New Directions

The current trend in Computer Vision and Natural Language Processing revolves around scaling, both in number of parameters and volume of data, using billion parameter methods trained over almost all accessible Internet content. This trend is exemplified by the Large Language Models (LLM) GPT-3 [22] or T5 [166] from the language side, ViT [44] and its variants from the vision side [11, 206, 207] and also the multimodal Vision and Language CLIP [164] and GPT-4 [156], which are trained by joint representations with contrastive loss from hundreds of millions of image-text pairs. Moreover, many works [32, 74, 75, 101, 116, 120, 133, 192, 197, 200, 208, 234, 240] have later combined the success of previous other LLM to create new variants and perform multiple tasks at the same time. However, LLM require such amount of computational resources that makes them infeasible for certain situations, like embedded systems or devices in a distributed Internet of Things setup. More importantly, not all applications dispose of vast amounts of available data. Therefore, the ability to adapt to new concepts or tasks with limited data, not only reduces the cost of collecting new data, but may be a necessary requirement for the given task. For instance, detecting and understanding rare diseases, or training a VQA system on an under-represented language or domain, may dispose of

limited available data. In such cases, an AI capable of discovering meaningful patterns within the limited data would greatly benefit the society.

The recent advancements and widespread adoption of AI in real-world applications, coupled with the increased accessibility to AI technologies for the general public, have raised significant concerns within society and authorities regarding the protection of privacy for sensitive data. This, linked to the aforementioned increased use of data to train the models, underscores the critical need for robust regulations and ethical frameworks to safeguard sensitive information and ensure responsible AI development. An information-retrieval system arranged as a VQA method following the paradigm of using NLP as interface, that operates in a real environment may need to process confidential medical records, financial information, personally identifiable data, as well as critical industry information of a company such as prices, stock, or customers details, which are susceptible to unintentional disclosure, data breaches, unauthorized access, and also the potential target for malicious attacks that exploit the data used to train AI methods. This growing concern emphasizes the need for privacy preserving methods [46], and urge the differential privacy community to move from toy and synthetic problems to more realistic ones, while the rest of the community, need to start integrating such algorithms, as safeguarding sensitive data will become an imperative requirement, and its importance will only continue to grow in the near future.

Furthermore, there has been a growing demand over the past few years to attain a deeper understanding of ML methods, which has significantly increased with the adoption of Transformers as the main backbone [211]. We firmly believe that explainability and interpretability must be regarded as critical aspects in future research, further amplified by the aforementioned widespread use of AI in real-world applications, and particularly in models deployed in critical scenarios. While we start addressing this problem through the answer page prediction sub-task in Chapter 8, there is still a lot of work to do in this direction to build more robust and interpretable methods.

Finally, the remarkable capability of recent methods to generate highly realistic text and images presents challenges in discerning between genuine and artificially generated content, which can contribute to the dissemination of misinformation or malicious activities. These AI-generated images will soon reach document field also, creating hard to identify forgeries or completely new documents, with potential malicious uses such as asking for fraudulent refunds, fake contracts or even fake signed certificates or letters that may make appear a person guilty of a crime. Thus, with the AI document information extraction systems such as VQA, realistic generated images and documents will pose new threats that will need to be tackled. For this, there is an urgent need for collaboration among various stakeholders within the community to engage in discussions regarding the ethical implications, conduct further research to deepen our understanding of ML methods, and implement appropriate regulations and security measures to address these concerns.

List of Contributions

Topics

The main topic of this dissertation is the incorporation of recognized text into the VQA task for natural scenes and documents, with the objective to create new datasets and methods that can read, reason and integrate the text with other visual cues. Nonetheless, this thesis has also led to additional collaborations on the same and other subjects which we list below:

International Journals

- **Real-time lexicon-free scene text retrieval.** Andrés Mafla, **Rubèn Tito**, Sounak Dey, Lluís Gómez, Marçal Rusiñol, Ernest Valveny, Dimosthenis Karatzas. *Pattern Recognition*, 2021.
- **Multimodal grid features and cell pointers for scene text visual question answering.** Lluís Gómez, Ali Furkan Biten, **Rubèn Tito**, Andrés Mafla, Marçal Rusiñol, Ernest Valveny, Dimosthenis Karatzas. *Pattern Recognition Letters*, 2021.
- **Hierarchical multimodal transformers for Multipage DocVQA.** **Rubèn Tito**, Dimosthenis Karatzas, Ernest Valveny. *Pattern Recognition*, 2023.

International Conferences

- **Scene Text Visual Question Answering.** Ali Furkan Biten*, **Rubèn Tito***, Andrés Mafla*, Lluís Gómez, Marçal Rusiñol, Ernest Valveny, CV Jawahar, Dimosthenis Karatzas. *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- **Document Collection Visual Question Answering.** **Rubèn Tito**, Dimosthenis Karatzas, and Ernest Valveny. *International Conference on Document Analysis and Recognition (ICDAR)*, 2021.

- **InfographicsVQA.** Minesh Mathew, Viraj Bagal, **Rubèn Tito**, Dimosthenis Karatzas, Ernest Valveny, CV Jawahar. *Winter Application in Computer Vision (WACV)*, 2022.
- **OCR-IDL: OCR Annotations for industry document library dataset.** Ali Furkan Biten, **Rubèn Tito**, Lluís Gomez, Ernest Valveny, Dimosthenis Karatzas. *IEEE/CVF European Conference on Computer Vision (ECCV)*, 2023
- **Document Understanding Dataset and Evaluation (DUDE).** Jordy Van Landeghem, **Rubèn Tito**, Łukasz Borchmann, Michal Pietruszka, Paweł Józiaś, Rafał Powalski, Dawid Jurkiewicz, Mickaël Coustaty, Bertrand Ackaert, Ernest Valveny, Matthew Blaschko, Sien Moens, Tomasz Stanisławek. *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.
- **[Competition]: ICDAR 2019 competition on scene text visual question answering.** Ali Furkan Biten*, **Rubèn Tito***, Andrés Mafla*, Lluís Gomez, Marçal Rusinol, Minesh Mathew, CV Jawahar, Ernest Valveny, Dimosthenis Karatzas. *International Conference on Document Analysis and Recognition (ICDAR)*, 2019.
- **[Competition]: Document visual question answering challenge 2020.** Minesh Mathew, **Rubèn Tito**, Dimosthenis Karatzas, R Manmatha, CV Jawahar. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Text and Documents in the Deep Learning Era workshop*, 2020.
- **[Competition]: ICDAR 2021 competition on document visual question answering.** **Rubèn Tito**, Minesh Mathew, CV Jawahar, Ernest Valveny, Dimosthenis Karatzas. *International Conference on Document Analysis and Recognition (ICDAR)*, 2021.
- **[Competition]: ICDAR 2023 Competition on Document Understanding of Everything (DUDE).** Jordy Van Landeghem, **Rubèn Tito**, Łukasz Borchmann, Michal Pietruszka, Dawid Jurkiewicz, Rafał Powalski, Paweł Józiaś, Sanket Biswas, Mickaël Coustaty, Tomasz Stanisławek. *International Conference on Document Analysis and Recognition (ICDAR)*, 2023.

Bibliography

- [1] Manoj Acharya, Kushal Kafle, and Christopher Kanan. Tallyqa: Answering complex counting questions. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 8076–8084, 2019.
- [2] Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Aniruddha Kembhavi. Don’t just assume; look and answer: Overcoming priors for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4971–4980, 2018.
- [3] Jon Almazán, Albert Gordo, Alicia Fornés, and Ernest Valveny. Word spotting and recognition with embedded attributes. *IEEE transactions on pattern analysis and machine intelligence*, 36(12):2552–2566, 2014.
- [4] Amazon. Amazon textract, 2021. Accessed on Aug. 31, 2023.
- [5] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6077–6086, 2018.
- [6] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015.
- [7] Srikar Appalaraju, Bhavan Jasani, Bhargava Urala Kota, Yusheng Xie, and R Manmatha. Docformer: End-to-end transformer for document understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 993–1003, 2021.
- [8] Youngmin Baek, Bado Lee, Dongyoon Han, Sangdoo Yun, and Hwalsuk Lee. Character region awareness for text detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9365–9374, 2019.
- [9] Fan Bai, Zhanzhan Cheng, Yi Niu, Shiliang Pu, and Shuigeng Zhou. Edit probability for scene text recognition. In *proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1508–1516, 2018.

- [10] Ankan Bansal, Yuting Zhang, and Rama Chellappa. Visual question answering on image sets. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pages 51–67. Springer, 2020.
- [11] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. In *ICLR 2022 - 10th International Conference on Learning Representations*, 2022.
- [12] Darwin Bautista and Rowel Atienza. Scene text recognition with permuted autoregressive sequence models. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVIII*, pages 178–196. Springer, 2022.
- [13] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *European conference on computer vision*, pages 404–417. Springer, 2006.
- [14] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- [15] Ali Furkan Biten, Ron Litman, Yusheng Xie, Srikar Appalaraju, and R Manmatha. Latr: Layout-aware transformer for scene-text vqa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16548–16558, 2022.
- [16] Ali Furkan Biten, Rubèn Tito, Lluís Gomez, Ernest Valveny, and Dimosthenis Karatzas. Ocr-idl: Ocr annotations for industry document library dataset. In *Computer Vision–ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IV*, pages 241–252. Springer, 2023.
- [17] Ali Furkan Biten, Rubèn Tito, Andres Mafla, Lluís Gomez, Marçal Rusinol, Minesh Mathew, CV Jawahar, Ernest Valveny, and Dimosthenis Karatzas. Icdar 2019 competition on scene text visual question answering. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 1563–1570. IEEE, 2019.
- [18] Ali Furkan Biten, Rubèn Tito, Andres Mafla, Lluís Gomez, Marçal Rusinol, Ernest Valveny, CV Jawahar, and Dimosthenis Karatzas. Scene text visual question answering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4291–4301, 2019.
- [19] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.
- [20] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146, 2017.

- [21] Fedor Borisjuk, Albert Gordo, and Viswanath Sivakumar. Rosetta: Large scale system for text detection and recognition in images. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 71–79, 2018.
- [22] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [23] Michal Busta, Lukas Neumann, and Jiri Matas. Deep textspotter: An end-to-end trainable scene text localization and recognition framework. In *Proceedings of the IEEE international conference on computer vision*, pages 2204–2212, 2017.
- [24] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020.
- [25] Prithvijit Chattopadhyay, Ramakrishna Vedantam, Ramprasaath R Selvaraju, Dhruv Batra, and Devi Parikh. Counting everyday objects in everyday scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1135–1144, 2017.
- [26] Ritwick Chaudhry, Sumit Shekhar, Utkarsh Gupta, Pranav Maneriker, Prann Bansal, and Ajay Joshi. Leaf-qa: Locate, encode & attend for figure question answering. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3512–3521, 2020.
- [27] Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. Reading wikipedia to answer open-domain questions. In *ACL 2017 - 55th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers)*, 2017.
- [28] Xi Chen, Josip Djolonga, Piotr Padlewski, Basil Mustafa, Soravit Changpinyo, Jialin Wu, Carlos Riquelme Ruiz, Sebastian Goodman, Xiao Wang, Yi Tay, et al. Pali-x: On scaling up a multilingual vision and language model. *arXiv preprint arXiv:2305.18565*, 2023.
- [29] Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. Pali: A jointly-scaled multilingual language-image model. *arXiv preprint arXiv:2209.06794*, 2022.
- [30] Xiaoxue Chen, Lianwen Jin, Yuanzhi Zhu, Canjie Luo, and Tianwei Wang. Text recognition in the wild: A survey. *ACM Computing Surveys (CSUR)*, 54(2):1–35, 2021.

- [31] Zhanzhan Cheng, Fan Bai, Yunlu Xu, Gang Zheng, Shiliang Pu, and Shuigeng Zhou. Focusing attention: Towards accurate text recognition in natural images. In *Proceedings of the IEEE international conference on computer vision*, pages 5076–5084, 2017.
- [32] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, March 2023.
- [33] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *EMNLP 2014 - 2014 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*, 2014.
- [34] Noam Chomsky. *Modular Approaches to the Study of the Mind*, volume 1. San Diego State Univ, 1984.
- [35] Chee-Kheng Ch’ng, Chee Seng Chan, and Cheng-Lin Liu. Total-text: toward orientation robustness in scene text detection. *International Journal on Document Analysis and Recognition (IJDAR)*, 23(1):31–52, 2020.
- [36] Bertrand Couasnon and Aurélie Lemaitre. Recognition of tables and forms, 2014.
- [37] Cheng Da, Peng Wang, and Cong Yao. Levenshtein ocr. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVIII*, pages 322–338. Springer, 2022.
- [38] Brian Davis, Bryan Morse, Brian Price, Chris Tensmeyer, Curtis Wigington, and Vlad Morariu. End-to-end document recognition and understanding with dessurt. In *Computer Vision–ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IV*, pages 280–296. Springer, 2023.
- [39] Dan Deng, Haifeng Liu, Xuelong Li, and Deng Cai. Pixellink: Detecting scene text via instance segmentation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- [40] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [41] Andreas R Dengel and Bertin Klein. smartfix: A requirements-driven system for document analysis and understanding. In *International Workshop on Document Analysis Systems*, pages 433–444. Springer, 2002.
- [42] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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, 2019.
- [43] David Doermann, Karl Tombre, et al. *Handbook of document image processing and recognition*, volume 1. Springer, 2014.
 - [44] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR 2021 - 9th International Conference on Learning Representations*, 2021.
 - [45] Matthew Dunn, Levent Sagun, Mike Higgins, V Ugur Guney, Volkan Cirik, and Kyunghyun Cho. Searchqa: A new q&a dataset augmented with context from a search engine. *arXiv preprint arXiv:1704.05179*, 2017.
 - [46] Cynthia Dwork. Differential privacy: A survey of results. In *Theory and Applications of Models of Computation: 5th International Conference, TAMC 2008, Xi'an, China, April 25-29, 2008. Proceedings 5*, pages 1–19. Springer, 2008.
 - [47] Shancheng Fang, Hongtao Xie, Yuxin Wang, Zhendong Mao, and Yongdong Zhang. Read like humans: Autonomous, bidirectional and iterative language modeling for scene text recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7098–7107, 2021.
 - [48] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. 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, 2016.
 - [49] Chenyu Gao, Qi Zhu, Peng Wang, Hui Li, Yuliang Liu, Anton Van den Hengel, and Qi Wu. Structured multimodal attentions for textvqa. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):9603–9614, 2021.
 - [50] Haoyuan Gao, Junhua Mao, Jie Zhou, Zhiheng Huang, Lei Wang, and Wei Xu. Are you talking to a machine? dataset and methods for multilingual image question. *Advances in neural information processing systems*, 28, 2015.
 - [51] Lukasz Garncarek, Rafal Powalski, Tomasz Stanislawek, Bartosz Topolski, Piotr Halama, Michal Turski, and Filip Graliński. Lambert: layout-aware language modeling for information extraction. In *International Conference on Document Analysis and Recognition*, pages 532–547. Springer, 2021.
 - [52] Felix A Gers, Jürgen Schmidhuber, and Fred Cummins. Learning to forget: Continual prediction with lstm. *Neural Computation*, 12(10):2451–2471, 2000.
 - [53] Lluís Gómez, Andrés Mafla, Marçal Rusinol, and Dimosthenis Karatzas. Single shot scene text retrieval. In *Proceedings of the European conference on computer vision (ECCV)*, pages 700–715, 2018.

- [54] Lluís Gomez, Yash Patel, Marçal Rusinol, Dimosthenis Karatzas, and CV Jawahar. Self-supervised learning of visual features through embedding images into text topic spaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4230–4239, 2017.
- [55] Raul Gomez, Baoguang Shi, Lluís Gomez, Lukas Numann, Andreas Veit, Jiri Matas, Serge Belongie, and Dimosthenis Karatzas. Icdar2017 robust reading challenge on coco-text. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, volume 1, pages 1435–1443. IEEE, 2017.
- [56] Google. Google ocr, 2020. Accessed on Aug. 31, 2023.
- [57] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017.
- [58] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, pages 369–376, 2006.
- [59] Caglar Gulcehre, Sungjin Ahn, Ramesh Nallapati, Bowen Zhou, and Yoshua Bengio. Pointing the unknown words. In *54th Annual Meeting of the Association for Computational Linguistics, ACL 2016 - Long Papers*, 2016.
- [60] Danna Gurari, Qing Li, Abigale J Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham. Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3608–3617, 2018.
- [61] Wei Han, Hantao Huang, and Tao Han. Finding the evidence: Localization-aware answer prediction for text visual question answering. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3118–3131, 2020.
- [62] Adam W Harley, Alex Ufkes, and Konstantinos G Derpanis. Evaluation of deep convolutional nets for document image classification and retrieval. In *2015 13th International Conference on Document Analysis and Recognition (ICDAR)*, pages 991–995. IEEE, 2015.
- [63] Guiqing He, Feng Li, Qiyao Wang, Zongwen Bai, and Yuelei Xu. A hierarchical sampling based triplet network for fine-grained image classification. *Pattern Recognition*, 115:107889, 2021.
- [64] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.

- [65] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [66] Pan He, Weilin Huang, Yu Qiao, Chen Loy, and Xiaoou Tang. Reading scene text in deep convolutional sequences. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.
- [67] Tong He, Zhi Tian, Weilin Huang, Chunhua Shen, Yu Qiao, and Changming Sun. An end-to-end textspotter with explicit alignment and attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5020–5029, 2018.
- [68] Himanshu. Detectron2 for document layout analysis. <https://github.com/hpanwar08/detectron2>, 2020.
- [69] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [70] Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Boyd Adriane. "spacy: Industrial-strength natural language processing in python", 2020.
- [71] Ronghang Hu, Amanpreet Singh, Trevor Darrell, and Marcus Rohrbach. Iterative answer prediction with pointer-augmented multimodal transformers for textvqa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [72] Yupan Huang, Tengchao Lv, Lei Cui, Yutong Lu, and Furu Wei. Layoutlmv3: Pre-training for document ai with unified text and image masking. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4083–4091, 2022.
- [73] Zheng Huang, Kai Chen, Jianhua He, Xiang Bai, Dimosthenis Karatzas, Shijian Lu, and CV Jawahar. Icdar2019 competition on scanned receipt ocr and information extraction. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 1516–1520. IEEE, 2019.
- [74] Zhicheng Huang, Zhaoyang Zeng, Yupan Huang, Bei Liu, Dongmei Fu, and Jianlong Fu. Seeing out of the box: End-to-end pre-training for vision-language representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12976–12985, 2021.
- [75] Zhicheng Huang, Zhaoyang Zeng, Bei Liu, Dongmei Fu, and Jianlong Fu. Pixelbert: Aligning image pixels with text by deep multi-modal transformers. *arXiv preprint arXiv:2004.00849*, 2020.
- [76] Drew A. Hudson and Christopher D. Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF Conference on CVPR*, volume 3 (8), 2019.

- [77] Jonathan J. Hull. A database for handwritten text recognition research. *IEEE Transactions on pattern analysis and machine intelligence*, 16(5):550–554, 1994.
- [78] Suzana Ilić, Edison Marrese-Taylor, Jorge Balazs, and Yutaka Matsuo. Deep contextualized word representations for detecting sarcasm and irony. In *Proceedings of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, 2018.
- [79] Donato Impedovo and Giuseppe Pirlo. Automatic signature verification: The state of the art. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(5):609–635, 2008.
- [80] Brian Kenji Iwana, Syed Tahseen Raza Rizvi, Sheraz Ahmed, Andreas Dengel, and Seiichi Uchida. Judging a book by its cover. *arXiv preprint arXiv:1610.09204*, 2016.
- [81] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Reading text in the wild with convolutional neural networks. *International journal of computer vision*, 116(1):1–20, 2016.
- [82] Tanuj Jain, Christopher Lennan, Zubin John, and Dat Tran. Imagededup. <https://github.com/ideal0/imagededup>, 2019.
- [83] Guillaume Jaume, Hazim Kemal Ekenel, and Jean-Philippe Thiran. Funsd: A dataset for form understanding in noisy scanned documents. In *2019 International Conference on Document Analysis and Recognition Workshops (ICDARW)*, volume 2, pages 1–6. IEEE, 2019.
- [84] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021.
- [85] Huaizu Jiang, Ishan Misra, Marcus Rohrbach, Erik Learned-Miller, and Xinlei Chen. In defense of grid features for visual question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10267–10276, 2020.
- [86] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2901–2910, 2017.
- [87] Karen Sparck Jones. A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation*, 1972.
- [88] Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *ACL 2017 - 55th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers)*, 2017.

- [89] Kushal Kafle and Christopher Kanan. Answer-type prediction for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4976–4984, 2016.
- [90] Kushal Kafle, Brian Price, Scott Cohen, and Christopher Kanan. Dvqa: Understanding data visualizations via question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5648–5656, 2018.
- [91] Samira Ebrahimi Kahou, Vincent Michalski, Adam Atkinson, Ákos Kádár, Adam Trischler, and Yoshua Bengio. Figureqa: An annotated figure dataset for visual reasoning. In *6th International Conference on Learning Representations, ICLR 2018 - Workshop Track Proceedings*, 2018.
- [92] Yash Kant, Dhruv Batra, Peter Anderson, Alexander Schwing, Devi Parikh, Jiasen Lu, and Harsh Agrawal. Spatially aware multimodal transformers for textvqa. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX 16*, pages 715–732. Springer, 2020.
- [93] Dimosthenis Karatzas, Lluís Gomez-Bigorda, Angelos Nicolaou, Suman Ghosh, Andrew Bagdanov, Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vijay Ramaseshan Chandrasekhar, Shijian Lu, et al. Icdar 2015 competition on robust reading. In *2015 13th international conference on document analysis and recognition (ICDAR)*, pages 1156–1160. IEEE, 2015.
- [94] Dimosthenis Karatzas, Faisal Shafait, Seiichi Uchida, Masakazu Iwamura, Lluís Gomez i Bigorda, Sergi Robles Mestre, Joan Mas, David Fernandez Mota, Jon Almazan Almazan, and Lluís Pere De Las Heras. Icdar 2013 robust reading competition. In *2013 12th international conference on document analysis and recognition*, pages 1484–1493. IEEE, 2013.
- [95] Isaak Kavasidis, Carmelo Pino, Simone Palazzo, Francesco Rundo, Daniela Giordano, P Messina, and Concetto Spampinato. A saliency-based convolutional neural network for table and chart detection in digitized documents. In *International conference on image analysis and processing*, pages 292–302. Springer, 2019.
- [96] Vahid Kazemi and Ali Elqursh. Show, ask, attend, and answer: A strong baseline for visual question answering. *arXiv preprint arXiv:1704.03162*, 2017.
- [97] Aniruddha Kembhavi, Minjoon Seo, Dustin Schwenk, Jonghyun Choi, Ali Farhadi, and Hannaneh Hajishirzi. Are you smarter than a sixth grader? textbook question answering for multimodal machine comprehension. In *Proceedings of the IEEE Conference on Computer Vision and Pattern recognition*, pages 4999–5007, 2017.
- [98] Aditya Khosla, Tinghui Zhou, Tomasz Malisiewicz, Alexei A Efros, and Antonio Torralba. Undoing the damage of dataset bias. In *European Conference on Computer Vision*, pages 158–171. Springer, 2012.

- [99] Jihyung Kil, Soravit Changpinyo, Xi Chen, Hexiang Hu, Sebastian Goodman, Wei-Lun Chao, and Radu Soricut. Prestu: Pre-training for scene-text understanding. *arXiv preprint arXiv:2209.05534*, 2022.
- [100] Geewook Kim, Teakgyu Hong, Moonbin Yim, JeongYeon Nam, Jinyoung Park, Jinyeong Yim, Wonseok Hwang, Sangdoo Yun, Dongyoon Han, and Seunghyun Park. Ocr-free document understanding transformer. In *European Conference on Computer Vision*, pages 498–517. Springer, 2022.
- [101] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *International Conference on Machine Learning*, pages 5583–5594. PMLR, 2021.
- [102] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, 2017.
- [103] Ivan Krasin, Tom Duerig, Neil Alldrin, Vittorio Ferrari, Sami Abu-El-Haija, Alina Kuznetsova, Hassan Rom, Jasper Uijlings, Stefan Popov, Shahab Kamali, Matteo Mallocci, Jordi Pont-Tuset, Andreas Veit, Serge Belongie, Victor Gomes, Abhinav Gupta, Chen Sun, Gal Chechik, David Cai, Zheyun Feng, Dhyanesh Narayanan, and Kevin Murphy. Openimages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from <https://storage.googleapis.com/openimages/web/index.html>*, 2(3):18, 2017.
- [104] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017.
- [105] Praveen Krishnan, Kartik Dutta, and CV Jawahar. Word spotting and recognition using deep embedding. In *2018 13th IAPR International Workshop on Document Analysis Systems (DAS)*, pages 1–6. IEEE, 2018.
- [106] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.
- [107] Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466, 2019.
- [108] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. In *8th International Conference on Learning Representations, ICLR 2020*, 2020.

- [109] Thomas K Landauer, Peter W Foltz, and Darrell Laham. An introduction to latent semantic analysis. *Discourse processes*, 25(2-3):259–284, 1998.
- [110] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [111] Chen-Yu Lee and Simon Osindero. Recursive recurrent nets with attention modeling for ocr in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2231–2239, 2016.
- [112] Kenton Lee, Mandar Joshi, Iulia Raluca Turc, Hexiang Hu, Fangyu Liu, Julian Martin Eisenschlos, Urvashi Khandelwal, Peter Shaw, Ming-Wei Chang, and Kristina Toutanova. Pix2struct: Screenshot parsing as pretraining for visual language understanding. In *International Conference on Machine Learning*, pages 18893–18912. PMLR, 2023.
- [113] Tao Lei, Yu Zhang, Sida I. Wang, Hui Dai, and Yoav Artzi. Simple recurrent units for highly parallelizable recurrence. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018*, 2018.
- [114] Vladimir I Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, pages 707–710. Soviet Union, 1966.
- [115] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 2020.
- [116] Gen Li, Nan Duan, Yuejian Fang, Ming Gong, and Daxin Jiang. Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34 (07), pages 11336–11344, 2020.
- [117] Junlong Li, Yiheng Xu, Tengchao Lv, Lei Cui, Cha Zhang, and Furu Wei. Dit: Self-supervised pre-training for document image transformer. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 3530–3539, 2022.
- [118] Minghao Li, Tengchao Lv, Jingye Chen, Lei Cui, Yijuan Lu, Dinei Florencio, Cha Zhang, Zhoujun Li, and Furu Wei. Trocr: Transformer-based optical character recognition with pre-trained models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37, 2023.
- [119] Peizhao Li, Jiuxiang Gu, Jason Kuen, Vlad I Morariu, Handong Zhao, Rajiv Jain, Varun Manjunatha, and Hongfu Liu. Selfdoc: Self-supervised document representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5652–5660, 2021.

- [120] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX 16*, pages 121–137. Springer, 2020.
- [121] Minghui Liao, Baoguang Shi, and Xiang Bai. Textboxes++: A single-shot oriented scene text detector. *IEEE transactions on image processing*, 27(8):3676–3690, 2018.
- [122] Minghui Liao, Baoguang Shi, Xiang Bai, Xinggang Wang, and Wenyu Liu. Textboxes: A fast text detector with a single deep neural network. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31, 2017.
- [123] Minghui Liao, Zhen Zhu, Baoguang Shi, Gui-song Xia, and Xiang Bai. Rotation-sensitive regression for oriented scene text detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5909–5918, 2018.
- [124] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [125] Tie-Yan Liu. Learning to rank for information retrieval. *Foundations and Trends in Information Retrieval*, 3(3):225–331, 2009.
- [126] Xiaojing Liu, Feiyu Gao, Qiong Zhang, and Huasha Zhao. Graph convolution for multimodal information extraction from visually rich documents. In *NAACL HLT 2019 - 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference*, volume 2, 2019.
- [127] Xuebo Liu, Ding Liang, Shi Yan, Dagui Chen, Yu Qiao, and Junjie Yan. Fots: Fast oriented text spotting with a unified network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5676–5685, 2018.
- [128] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [129] Yuliang Liu, Lianwen Jin, Shuaitao Zhang, Canjie Luo, and Sheng Zhang. Curved scene text detection via transverse and longitudinal sequence connection. *Pattern Recognition*, 90:337–345, 2019.
- [130] Zichuan Liu, Yixing Li, Fengbo Ren, Wang Ling Goh, and Hao Yu. Squeezedtext: A real-time scene text recognition by binary convolutional encoder-decoder network. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.

- [131] Shangbang Long, Xin He, and Cong Yao. Scene text detection and recognition: The deep learning era. *International Journal of Computer Vision*, 129:161–184, 2021.
- [132] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004.
- [133] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32, 2019.
- [134] Xiaopeng Lu, Zhen Fan, Yansen Wang, Jean Oh, and Carolyn P Rosé. Localize, group, and select: Boosting text-vqa by scene text modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2631–2639, 2021.
- [135] Pengyuan Lyu, Minghui Liao, Cong Yao, Wenhao Wu, and Xiang Bai. Mask textspotter: An end-to-end trainable neural network for spotting text with arbitrary shapes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 67–83, 2018.
- [136] Jianqi Ma, Weiyuan Shao, Hao Ye, Li Wang, Hong Wang, Yingbin Zheng, and Xiangyang Xue. Arbitrary-oriented scene text detection via rotation proposals. *IEEE transactions on multimedia*, 20(11):3111–3122, 2018.
- [137] Mateusz Malinowski and Mario Fritz. A multi-world approach to question answering about real-world scenes based on uncertain input. In *Advances in Neural Information Processing Systems*, volume 2, 2014.
- [138] Mateusz Malinowski, Marcus Rohrbach, and Mario Fritz. Ask your neurons: A neural-based approach to answering questions about images. In *Proceedings of the IEEE international conference on computer vision*, pages 1–9, 2015.
- [139] R Manmatha and WB Croft. Word spotting: Indexing handwritten archives. *Intelligent Multimedia Information Retrieval Collection*, pages 43–64, 1997.
- [140] Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. Okvqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition*, pages 3195–3204, 2019.
- [141] David Marr. *Vision: A computational investigation into the human representation and processing of visual information*. MIT press, 2010.
- [142] Minesh Mathew, Viraj Bagal, Rubèn Tito, Dimosthenis Karatzas, Ernest Valveny, and CV Jawahar. Infographicvqa. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1697–1706, 2022.
- [143] Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. Docvqa: A dataset for vqa on document images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2200–2209, 2021.

- [144] Minesh Mathew, Ruben Tito, Dimosthenis Karatzas, R Manmatha, and CV Jawahar. Document visual question answering challenge 2020. *arXiv preprint arXiv:2008.08899*, 2020.
- [145] Bryan McCann, James Bradbury, Caiming Xiong, and Richard Socher. Learned in translation: Contextualized word vectors. *Advances in neural information processing systems*, 30, 2017.
- [146] Microsoft. Azure, 2021. Accessed on Aug. 31, 2023.
- [147] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 26, 2013.
- [148] Anand Mishra, Karteek Alahari, and CV Jawahar. Scene text recognition using higher order language priors. In *BMVC-British machine vision conference*. BMVA, 2012.
- [149] Anand Mishra, Karteek Alahari, and CV Jawahar. Image retrieval using textual cues. In *Proceedings of the IEEE international conference on computer vision*, pages 3040–3047, 2013.
- [150] Anand Mishra, Shashank Shekhar, Ajeet Kumar Singh, and Anirban Chakraborty. Ocr-vqa: Visual question answering by reading text in images. In *2019 international conference on document analysis and recognition (ICDAR)*, pages 947–952. IEEE, 2019.
- [151] Harold Mouchère, Christian Viard-Gaudin, Richard Zanibbi, and Utpal Garain. Icfhr2016 crohme: Competition on recognition of online handwritten mathematical expressions. In *2016 15th International Conference on Frontiers in Handwriting Recognition (ICFHR)*, pages 607–612. IEEE, 2016.
- [152] Byeonghu Na, Yoonsik Kim, and Sungrae Park. Multi-modal text recognition networks: Interactive enhancements between visual and semantic features. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVIII*, pages 446–463. Springer, 2022.
- [153] Nibal Nayef, Yash Patel, Michal Busta, Pinaki Nath Chowdhury, Dimosthenis Karatzas, Wafa Khelif, Jiri Matas, Umapada Pal, Jean-Christophe Burie, Cheng-lin Liu, et al. Icdar2019 robust reading challenge on multi-lingual scene text detection and recognition—rrc-mlt-2019. In *2019 International conference on document analysis and recognition (ICDAR)*, pages 1582–1587. IEEE, 2019.
- [154] Neha. Movie genre from its poster. <https://www.kaggle.com/datasets/neh1703/movie-genre-from-its-poster/>, 2018.
- [155] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. Ms marco: A human generated machine reading comprehension dataset. In *CEUR Workshop Proceedings*, volume 1773, 2016.

- [156] OpenAI. Gpt-4 technical report, 2023.
- [157] Rasmus Berg Palm, Ole Winther, and Florian Laws. Cloudscan-a configuration-free invoice analysis system using recurrent neural networks. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, volume 1, pages 406–413. IEEE, 2017.
- [158] Yash Patel, Lluís Gomez, Marçal Rusinol, and Dimosthenis Karatzas. Dynamic lexicon generation for natural scene images. In *European Conference on Computer Vision*, pages 395–410. Springer, 2016.
- [159] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014.
- [160] Trung Quy Phan, Palaiahnakote Shivakumara, Shangxuan Tian, and Chew Lim Tan. Recognizing text with perspective distortion in natural scenes. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 569–576, 2013.
- [161] Rafal Powalski, Łukasz Borchmann, Dawid Jurkiewicz, Tomasz Dwojak, Michał Pietruszka, and Gabriela Palka. Going full-tilt boogie on document understanding with text-image-layout transformer. In *Document Analysis and Recognition—ICDAR 2021: 16th International Conference, Lausanne, Switzerland, September 5–10, 2021, Proceedings, Part II 16*, pages 732–747. Springer, 2021.
- [162] Le Qi, Shangwen Lv, Hongyu Li, Jing Liu, Yu Zhang, Qiaoqiao She, Hua Wu, Haifeng Wang, and Ting Liu. Dureadervis: A: A chinese dataset for open-domain document visual question answering. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1338–1351, 2022.
- [163] Siyang Qin, Alessandro Bissacco, Michalis Raptis, Yasuhisa Fujii, and Ying Xiao. Towards unconstrained end-to-end text spotting. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4704–4714, 2019.
- [164] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [165] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training, 2018.
- [166] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67, 2020.
- [167] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions for machine comprehension of text. In *EMNLP 2016 - Conference on Empirical Methods in Natural Language Processing, Proceedings*, 2016.

- [168] Toni M Rath and Raghavan Manmatha. Word image matching using dynamic time warping. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 2, pages II–II. IEEE, 2003.
- [169] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.
- [170] Mengye Ren, Ryan Kiros, and Richard S. Zemel. Exploring models and data for image question answering. In *Advances in Neural Information Processing Systems*, volume 2015-January, 2015.
- [171] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015.
- [172] Anhar Risnumawan, Palaiahankote Shivakumara, Chee Seng Chan, and Chew Lim Tan. A robust arbitrary text detection system for natural scene images. *Expert Systems with Applications*, 41(18):8027–8048, 2014.
- [173] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [174] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science, 1985.
- [175] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.
- [176] Daniel Schuster, Klemens Muthmann, Daniel Esser, Alexander Schill, Michael Berger, Christoph Weidling, Kamil Aliyev, and Andreas Hofmeier. Intellix–end-user trained information extraction for document archiving. In *2013 12th International Conference on Document Analysis and Recognition*, pages 101–105. IEEE, 2013.
- [177] Mike Schuster and Kuldip K Paliwal. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing*, 45(11):2673–2681, 1997.
- [178] 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, 2016.

- [179] Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hananneh Hajishirzi. Bidirectional attention flow for machine comprehension. In *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, 2017.
- [180] Baoguang Shi, Xiang Bai, and Serge Belongie. Detecting oriented text in natural images by linking segments. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2550–2558, 2017.
- [181] Baoguang Shi, Xiang Bai, and Cong Yao. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *IEEE transactions on pattern analysis and machine intelligence*, 39(11):2298–2304, 2016.
- [182] Baoguang Shi, Xinggang Wang, Pengyuan Lyu, Cong Yao, and Xiang Bai. Robust scene text recognition with automatic rectification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4168–4176, 2016.
- [183] Baoguang Shi, Mingkun Yang, Xinggang Wang, Pengyuan Lyu, Cong Yao, and Xiang Bai. Aster: An attentional scene text recognizer with flexible rectification. *IEEE transactions on pattern analysis and machine intelligence*, 41(9):2035–2048, 2018.
- [184] Kevin J Shih, Saurabh Singh, and Derek Hoiem. Where to look: Focus regions for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4613–4621, 2016.
- [185] Noah Siegel, Zachary Horvitz, Roie Levin, Santosh Divvala, and Ali Farhadi. Figureseer: Parsing result-figures in research papers. In *ECCV*, pages 664–680. Springer, 2016.
- [186] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, 2015.
- [187] Ajeet Kumar Singh, Anand Mishra, Shashank Shekhar, and Anirban Chakraborty. From strings to things: Knowledge-enabled vqa model that can read and reason. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4602–4612, 2019.
- [188] Amanpreet Singh, Vedanuj Goswami, Vivek Natarajan, Yu Jiang, Xinlei Chen, Meet Shah, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. Mmf: A multimodal framework for vision and language research, 2020.
- [189] Amanpreet Singh, Vedanuj Goswami, and Devi Parikh. Are we pretraining it right? digging deeper into visio-linguistic pretraining. *arXiv preprint arXiv:2004.08744*, 2020.

- [190] Amanpreet Singh, Vivek Natarajan, Yu Jiang, Xinlei Chen, Meet Shah, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. Pythia-a platform for vision & language research. In *SysML Workshop, NeurIPS*, volume 2018, 2018.
- [191] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE/CVF CVPR*, pages 8317–8326, 2019.
- [192] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vlb-ert: Pre-training of generic visual-linguistic representations. In *8th International Conference on Learning Representations, ICLR 2020*, 2020.
- [193] Sebastian Sudholt and Gernot A Fink. Evaluating word string embeddings and loss functions for cnn-based word spotting. In *2017 14th iapr international conference on document analysis and recognition (icdar)*, volume 1, pages 493–498. IEEE, 2017.
- [194] Yipeng Sun, Jiaming Liu, Wei Liu, Junyu Han, Errui Ding, and Jingtuo Liu. Chinese street view text: Large-scale chinese text reading with partially supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9086–9095, 2019.
- [195] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27, 2014.
- [196] Alon Talmor and Jonathan Berant. The web as a knowledge-base for answering complex questions. In *NAACL HLT 2018 - 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference*, 2018.
- [197] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. In *EMNLP-IJCNLP 2019 - 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing, Proceedings of the Conference*, 2019.
- [198] Ryota Tanaka, Kyosuke Nishida, and Sen Yoshida. Visualmrc: Machine reading comprehension on document images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13878–13888, 2021.
- [199] Zineng Tang, Ziyi Yang, Guoxin Wang, Yuwei Fang, Yang Liu, Chenguang Zhu, Michael Zeng, Cha Zhang, and Mohit Bansal. Unifying vision, text, and layout for universal document processing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19254–19264, 2023.
- [200] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. Alpaca: A strong, replicable instruction-following model. *Stanford Center for Research on Foundation Models*. <https://crfm.stanford.edu/2023/03/13/alpaca.html>, 3(6):7, 2023.

- [201] Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhagen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. Movieqa: Understanding stories in movies through question-answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4631–4640, 2016.
- [202] Joshua B Tenenbaum and William T Freeman. Separating style and content with bilinear models. *Neural computation*, 12(6):1247–1283, 2000.
- [203] Zhi Tian, Weilin Huang, Tong He, Pan He, and Yu Qiao. Detecting text in natural image with connectionist text proposal network. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*, pages 56–72. Springer, 2016.
- [204] Rubèn Tito, Dimosthenis Karatzas, and Ernest Valveny. Document collection visual question answering. In *International Conference on Document Analysis and Recognition*, pages 778–792. Springer, 2021.
- [205] Rubèn Tito, Minesh Mathew, CV Jawahar, Ernest Valveny, and Dimosthenis Karatzas. Icdar 2021 competition on document visual question answering. In *International Conference on Document Analysis and Recognition*, pages 635–649. Springer, 2021.
- [206] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International conference on machine learning*, pages 10347–10357. PMLR, 2021.
- [207] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 32–42, 2021.
- [208] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [209] Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordoni, Philip Bachman, and Kaheer Suleman. Newsqa: A machine comprehension dataset. In *Proceedings of the 2nd Workshop on Representation Learning for NLP, Rep4NLP 2017 at the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017*, 2017.
- [210] Alexander Trott, Caiming Xiong, and Richard Socher. Interpretable counting for visual question answering. In *6th International Conference on Learning Representations, ICLR 2018 - Conference Track Proceedings*, 2018.
- [211] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

- [212] Andreas Veit, Tomas Matera, Lukas Neumann, Jiri Matas, and Serge Belongie. Coco-text: Dataset and benchmark for text detection and recognition in natural images. *arXiv preprint arXiv:1601.07140*, 2016.
- [213] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. *Advances in neural information processing systems*, 28, 2015.
- [214] Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. Git: A generative image-to-text transformer for vision and language. *arXiv preprint arXiv:2205.14100*, 2022.
- [215] Jun Wang, Mingfei Gao, Yuqian Hu, Ramprasaath R Selvaraju, Chetan Ramaiah, Ran Xu, Joseph F JaJa, and Larry S Davis. Tag: Boosting text-vqa via text-aware visual question-answer generation. *arXiv preprint arXiv:2208.01813*, 2022.
- [216] Kai Wang, Boris Babenko, and Serge Belongie. End-to-end scene text recognition. In *2011 International conference on computer vision*, pages 1457–1464. IEEE, 2011.
- [217] Kai Wang and Serge Belongie. Word spotting in the wild. In *European conference on computer vision*, pages 591–604. Springer, 2010.
- [218] Peng Wang, Qi Wu, Chunhua Shen, Anthony Dick, and Anton van den Hengel. Explicit knowledge-based reasoning for visual question answering. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 1290–1296, 2017.
- [219] Peng Wang, Qi Wu, Chunhua Shen, Anthony Dick, and Anton Van Den Hengel. Fvqa: Fact-based visual question answering. *IEEE transactions on pattern analysis and machine intelligence*, 40(10):2413–2427, 2017.
- [220] Wenhai Wang, Enze Xie, Xiang Li, Wenbo Hou, Tong Lu, Gang Yu, and Shuai Shao. Shape robust text detection with progressive scale expansion network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9336–9345, 2019.
- [221] Xinyu Wang, Yuliang Liu, Chunhua Shen, Chun Chet Ng, Canjie Luo, Lianwen Jin, Chee Seng Chan, Anton van den Hengel, and Liangwei Wang. On the general value of evidence, and bilingual scene-text visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10126–10135, 2020.
- [222] Tomas Wilkinson, Jonas Lindstrom, and Anders Brun. Neural ctrl-f: segmentation-free query-by-string word spotting in handwritten manuscript collections. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4433–4442, 2017.

- [223] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, pages 38–45, 2020.
- [224] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016.
- [225] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- [226] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017.
- [227] Yang Xu, Yiheng Xu, Tengchao Lv, Lei Cui, Furu Wei, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, Wanxiang Che, et al. Layoutlmv2: Multi-modal pre-training for visually-rich document understanding. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2579–2591, 2021.
- [228] Yiheng Xu, Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, and Ming Zhou. Layoutlm: Pre-training of text and layout for document image understanding. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1192–1200, 2020.
- [229] Semih Yagcioglu, Aykut Erdem, Erkut Erdem, and Nazli Ikizler-Cinbis. Recipeqa: A challenge dataset for multimodal comprehension of cooking recipes. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018*, 2018.
- [230] Zhengyuan Yang, Yijuan Lu, Jianfeng Wang, Xi Yin, Dinei Florencio, Lijuan Wang, Cha Zhang, Lei Zhang, and Jiebo Luo. Tap: Text-aware pre-training for text-vqa and text-caption. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8751–8761, 2021.
- [231] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32, 2019.

- [232] Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov, and Christopher D Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018*, 2018.
- [233] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked attention networks for image question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 21–29, 2016.
- [234] Fei Yu, Jiji Tang, Weichong Yin, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. Ernie-vil: Knowledge enhanced vision-language representations through scene graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35 (04), pages 3208–3216, 2021.
- [235] Licheng Yu, Eunbyung Park, Alexander C Berg, and Tamara L Berg. Visual madlibs: Fill in the blank description generation and question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2461–2469, 2015.
- [236] Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. Big bird: Transformers for longer sequences. *Advances in Neural Information Processing Systems*, 33:17283–17297, 2020.
- [237] Xiaoyu Zeng, Yanan Wang, Tai-Yin Chiu, Nilavra Bhattacharya, and Danna Gurari. Vision skills needed to answer visual questions. *Proc. ACM Hum.-Comput. Interact.*, 4(CSCW2), October 2020.
- [238] Peng Zhang, Yash Goyal, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Yin and yang: Balancing and answering binary visual questions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5014–5022, 2016.
- [239] Peng Zhang, Yunlu Xu, Zhanzhan Cheng, Shiliang Pu, Jing Lu, Liang Qiao, Yi Niu, and Fei Wu. Trie: end-to-end text reading and information extraction for document understanding. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 1413–1422, 2020.
- [240] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. Vinyl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5579–5588, 2021.
- [241] Xiang Zhang, Yongwen Su, Subarna Tripathi, and Zhuowen Tu. Text spotting transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9519–9528, 2022.

- [242] Zheng Zhang, Chengquan Zhang, Wei Shen, Cong Yao, Wenyu Liu, and Xiang Bai. Multi-oriented text detection with fully convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4159–4167, 2016.
- [243] Shuai Zhao, Xiaohan Wang, Linchao Zhu, and Yi Yang. Clip4str: A simple baseline for scene text recognition with pre-trained vision-language model. *arXiv preprint arXiv:2305.14014*, 2023.
- [244] Xu Zhong, Jianbin Tang, and Antonio Jimeno Yepes. Publaynet: largest dataset ever for document layout analysis. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 1015–1022. IEEE, 2019.
- [245] Bolei Zhou, Yuandong Tian, Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. Simple baseline for visual question answering. *arXiv preprint arXiv:1512.02167*, 2015.
- [246] Xinyu Zhou, Cong Yao, He Wen, Yuzhi Wang, Shuchang Zhou, Weiran He, and Jiajun Liang. East: an efficient and accurate scene text detector. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 5551–5560, 2017.
- [247] Yu Zhou, Xiaoni Li, Yucan Zhou, Yu Wang, Qinghua Hu, and Weiping Wang. Deep collaborative multi-task network: A human decision process inspired model for hierarchical image classification. *Pattern Recognition*, 124:108449, 2022.
- [248] Fengbin Zhu, Wenqiang Lei, Fuli Feng, Chao Wang, Haozhou Zhang, and Tat-Seng Chua. Towards complex document understanding by discrete reasoning. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4857–4866, 2022.
- [249] Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. Visual7w: Grounded question answering in images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4995–5004, 2016.

