

# Learning to Extract Semantic Structure from Documents Using Multimodal Fully Convolutional Neural Networks

## SUPPLEMENTAL MATERIAL

Xiao Yang<sup>‡</sup>, Ersin Yumer<sup>†</sup>, Paul Asente<sup>†</sup>, Mike Kraley<sup>†</sup>, Daniel Kifer<sup>‡</sup>, C. Lee Giles<sup>‡</sup>

<sup>‡</sup>The Pennsylvania State University    <sup>†</sup>Adobe Research

xuy111@psu.edu {yumer, asente, mkraley}@adobe.com dkifer@cse.psu.edu giles@ist.psu.edu

### 1. Synthetic Document Data

We introduced two methods to generate documents. In the first method, we generate LaTeX source files in which elements like paragraphs, figures, tables, captions, section headings and lists are randomly arranged using the “textblock” environment from the “textpos” package. Compiling these LaTeX files gives single, double, or triple-column PDFs. The generation process is summarized in Algorithm 1.

---

**Algorithm 1** Synthetic Document Generation

---

- 1:  $s \leftarrow$  a string containing preamble and necessary packages of a LaTeX source file
- 2: Select a LaTeX source file type  $T \in \{\text{single-column, double-column, triple-column}\}$
- 3: **while** space remains on the page **do**
- 4:   Select an element type  $E \in \{\text{figure, table, caption, section heading, list, paragraph}\}$
- 5:   Select an example  $e$  of type  $E$
- 6:    $s_e \leftarrow$  a string of LaTeX code that generates  $e$  using the “textblock” environment
- 7:    $s \leftarrow s + s_e$
- 8: **end while**

**Output:**  $s$

**Output:** A PDF document after compiling  $s$

---

Elements in a document are carefully selected following the guidelines below. Figure 1 shows several examples of the figures and tables used in the synthetic data generation.

- Candidate figures include natural images from MS COCO [3], academic-style figures and graphic drawings downloaded using web image search.
- Candidates tables include table images downloaded using web image search. Various queries are used to increase the diversity of downloaded tables.

- For paragraphs, we randomly sample sentences from a 2016 English Wikipedia dump [2].
- For section headings, we sample sentences and phrases that are section or subsection headings in the “Contents” block in a Wikipedia page.
- For lists, we sample list items from Wikipedia pages, ensuring that all items in a list come from the same Wikipedia page.
- For captions, we either use the associated caption (for images from MS COCO) or the title of the image in web image search, which can be found in the span with class name “irc\_pt”.

In the second document generation method, we collected and labeled 271 documents with varied, complicated layouts. We then randomly replaced each element with a standalone paragraph, figure, table, caption, section heading or list generated as stated above. Figure 2 shows several examples from the 271 documents.

### 2. Visualizing the Segmentation Results

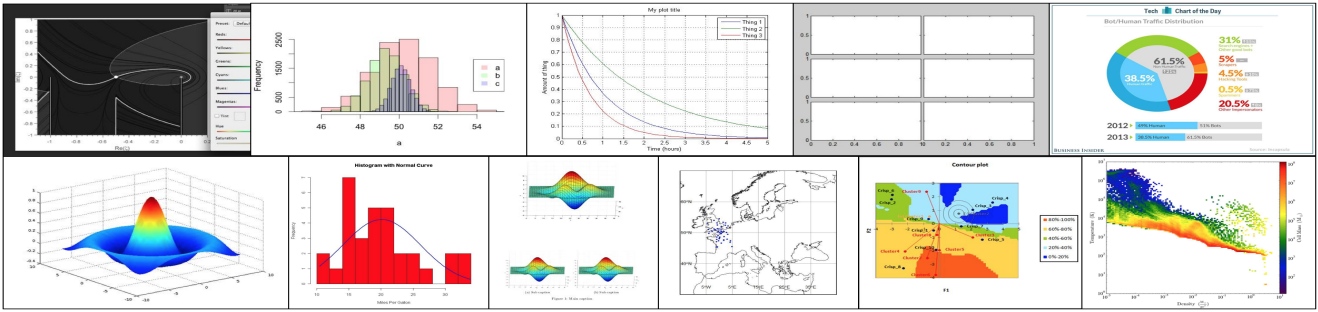
Each pixel  $p$  in the model’s output layer is assigned the color of the most likely class label  $l$ . The RGB value of that color is then weighted by the probability  $P(l)$ .

### 3. Post-processing


We apply an optional post-processing step to clean up segment masks for documents in PDF format. First, we obtain candidate bounding boxes by using the auto-tagging capabilities of Adobe Acrobat [1] and parsing the results. Boxes are stored in a tree structure, and each node’s box can be a TextRun (a sequence of characters), TextLine (potentially a text line), Paragraph (potentially a paragraph) or Container (potentially figures or tables). Note that we



(1)



(2)



(3)

Operating result	2H 2010	1H 2010	2010	2009	2008
Operating result	34	14	48	-7	-7
Restructuring costs	0	0	0	-21	0
Financial income and expenses	-2	-6	-6	0	0
Realised exchange rate differences UK	-10	0	-10	0	0
Share of profit of associates	1	0	1	0	0
<b>Result before tax</b>	<b>23</b>	<b>6</b>	<b>31</b>	<b>-46</b>	<b>0</b>
Income tax	-11	-4	-15	6	6
<b>Result after tax</b>	<b>12</b>	<b>4</b>	<b>16</b>	<b>-40</b>	<b>0</b>

Men's Standard	A	B	Women's Standard	A	B
2 Entries	22.35	22.90	1 Entry	25.43	26.96
100m Freestyle	49.23	50.44	500m Freestyle	55.24	56.60
200m Freestyle	1:48.72	1:51.40	1500m Freestyle	1:58.29	2:02.24
400m Freestyle	3:49.96	3:55.63	800m Freestyle	4:11.26	4:17.64
800m Freestyle	8:10.26	8:37.52	1500m Freestyle	8:35.98	8:48.70
1500m Freestyle	15:13.16	15:35.67	50m Backstroke	16:41.49	17:10.88
50m Backstroke	25.34	25.98	100m Backstroke	29.07	29.80
100m Backstroke	55.14	56.50	200m Backstroke	1:01.70	1:03.22
200m Backstroke	1:59.72	2:02.67		2:12.73	2:16.01

Equations (x = density)	y = 11.8	y = 37.7	y = 73.3	y = 124.0	y = 155.9	y = 235.39 - 171
5,000	13.5	41.3	72.9	115.4	142.9	162.8
7,000	11.1	38.0	71.3	115.5	139.6	161.8
9,000	10.7	37.4	70.9	121.1	154.7	180.8
11,000	11.5	39.7	76.6	127.9	157.6	184.1
13,000	12.3	38.8	76.6	129.8	162.3	189.3
15,000	10.7	36.6	73.7	127.3	158.0	189.9
17,000	12.2	36.4	72.9	131.7	170.4	205.9
19,000	12.2	29.9	54.9	103.6	140.9	175.8
21,000	12.2	41.5	84.9	140.3	176.9	207.2

(4)

Figure 1: Sample figures and tables used in synthetic documents generation. (1) Natural images from MS COCO dataset. (2) Academic-style figures from web image search. (3) Symbols and graphic drawings from web image search. (4) Tables from web image search.

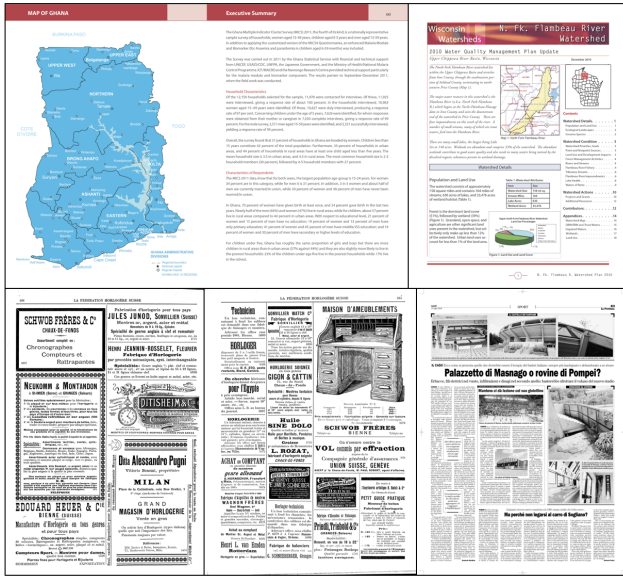


Figure 2: Examples of documents with complicated layout. We labeled regions in each document and then randomly replaced them with a standalone paragraph, figure, table, caption, section heading or list, as described in Sec. 1

ignore the semantic meanings associated with these boxes and only use the boxes as candidate bounding boxes in post-processing. Figure 3 (2) and 4 (2) illustrate candidate bounding boxes for each document.

---

### Algorithm 2 Segmentation Post-processing

---

**Input:**  $P \leftarrow$  probability map,  $P(u, v) \in R^{|C|}$  is a vector containing the probability of each class  $c \in C$  at location  $(u, v)$

**Input:**  $Boxes \leftarrow$  candidate bounding boxes

- 1:  $S \leftarrow$  segmentation to be generated
- 2: **for** each location  $(x, y) \in S$  **do**
- 3:      $S(x, y) \leftarrow$  background
- 4: **end for**
- 5: **for** each  $b \in Boxes$  **do**     ▷ parent box comes before child boxes
- 6:      $\bar{p} \leftarrow \sum_{(u,v) \in b} P(u, v)$
- 7:      $l \leftarrow \operatorname{argmax} \bar{p}$
- 8:     **for** each location  $(u, v) \in b$  **do**
- 9:         **if**  $S(u, v)$  is background **then**
- 10:              $S(u, v) \leftarrow l$
- 11:         **end if**
- 12:     **end for**
- 13: **end for**

**Output:**  $S$

---

Using these bounding box candidates, we refine the segmentation masks by first calculating the average class probability for pixels belonging to the same box, followed by assigning the most likely label to these pixels. The process

is summarized in Algorithm 2.

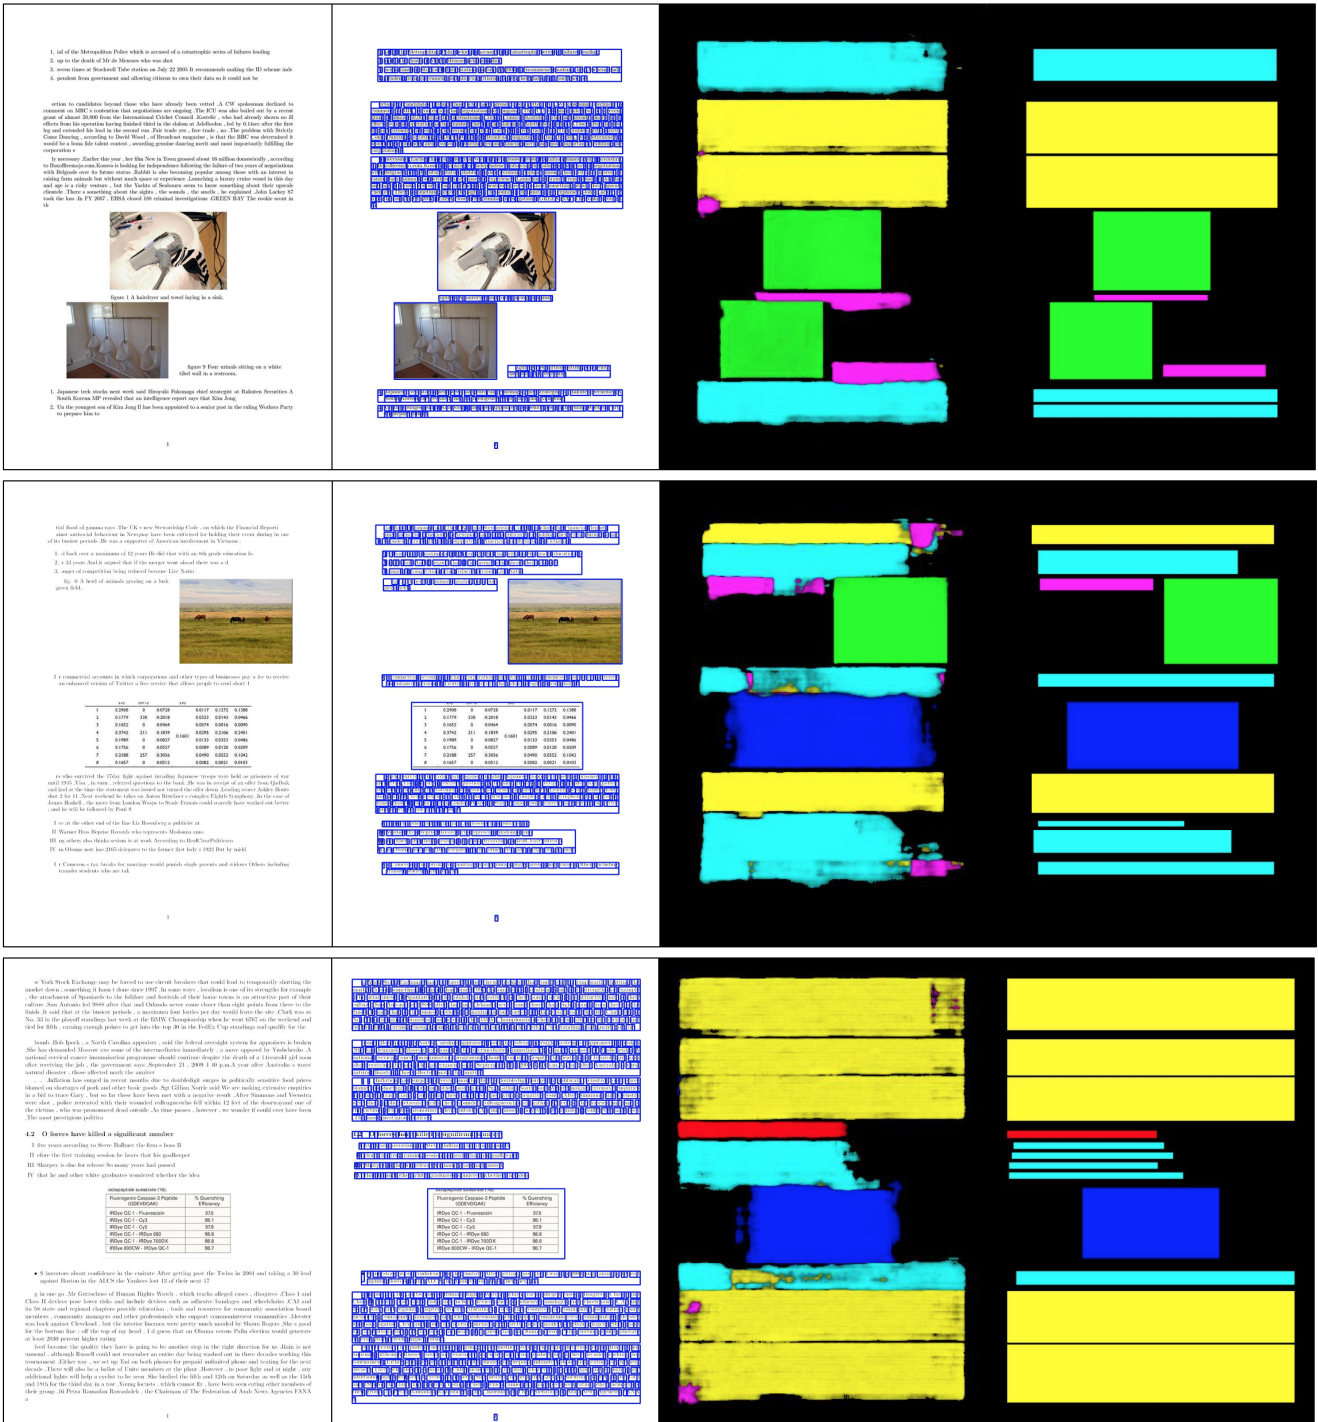
## 4. Additional Visualization Results

Figures 3 and 4 show additional visualization examples of synthetic documents, and Figure 5 shows additional examples of real documents.

## References

- [1] Adobe Acrobat. <http://www.adobe.com/accessibility/products/acrobat.html>. 1
- [2] Wikipedia. <https://dumps.wikimedia.org/>. 1
- [3] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pages 740–755. Springer, 2014. 1





(1)

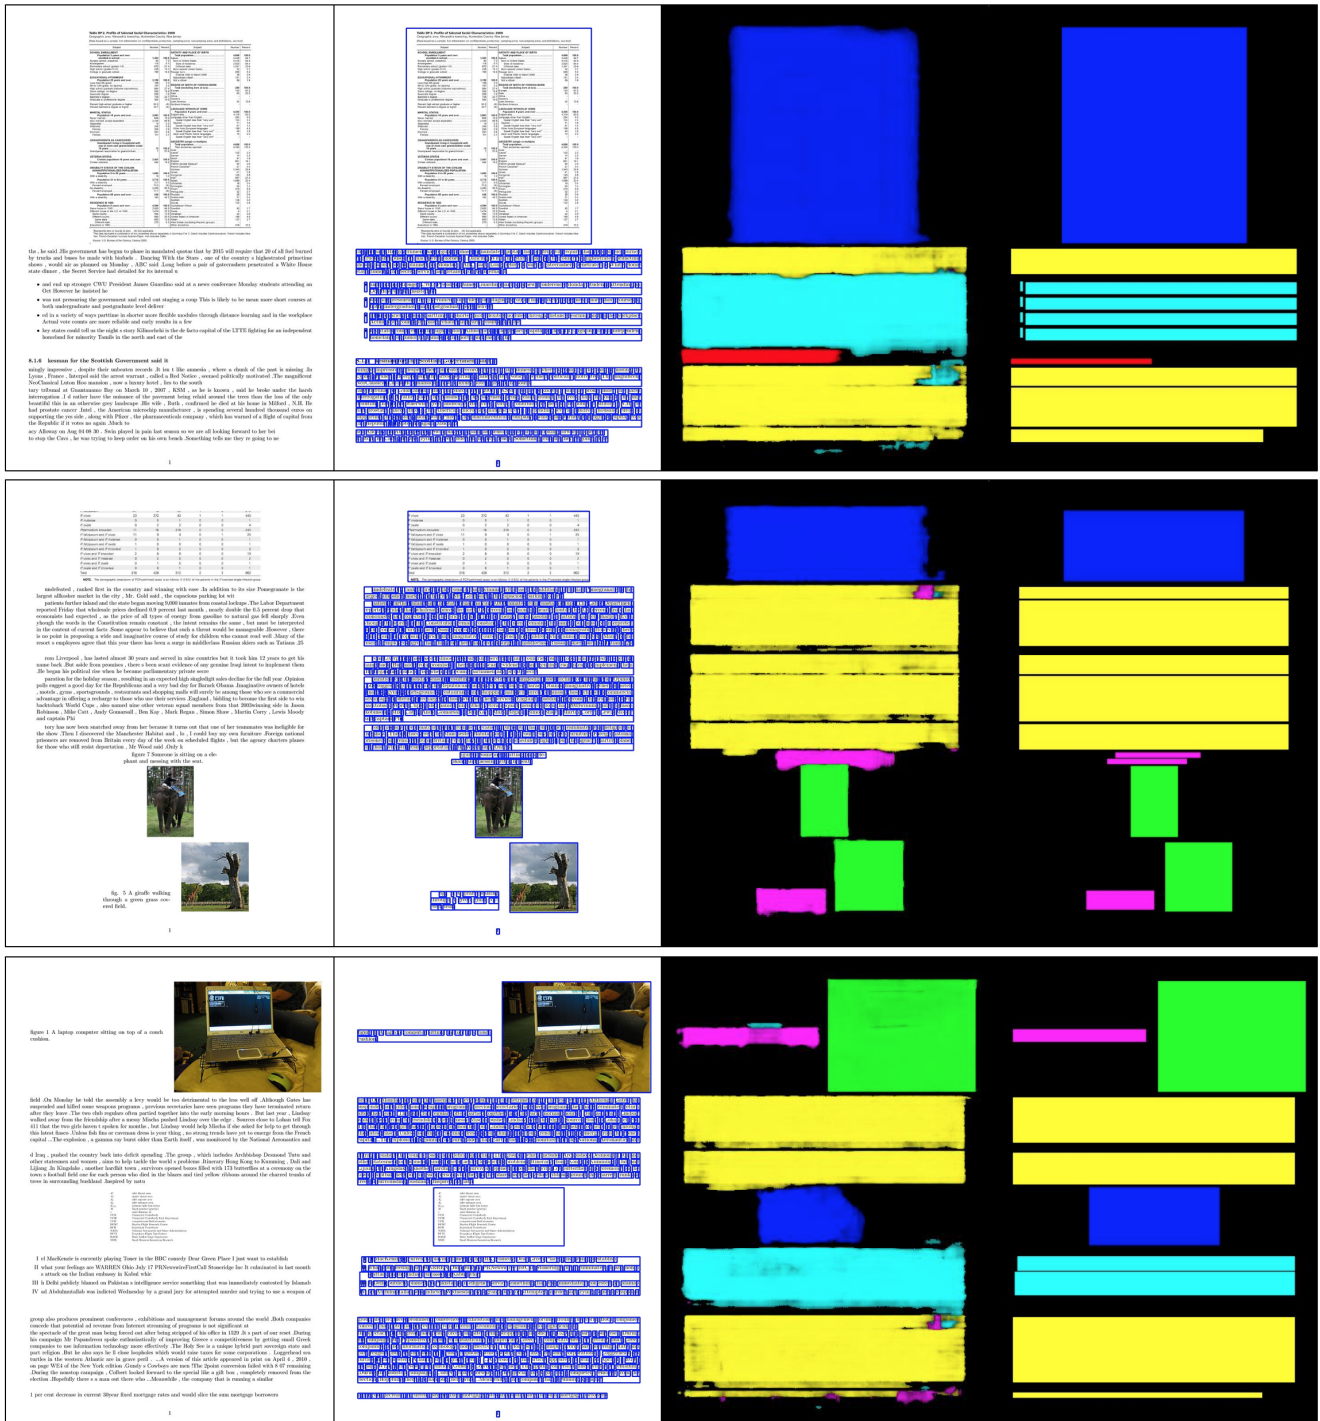
(2)

(3)

(4)

Figure 3: Synthetic documents and the corresponding segmentations. (1) Input synthetic documents. (2) Candidate bounding boxes obtained by parsing the PDF rendering commands. (3) Raw segmentation outputs. (4) Segmentations after post-processing. Segmentation label colors are: **figure**, **table**, **section heading**, **caption**, **list** and **paragraph**.





(1)

(2)

(3)

(4)

Figure 4: Synthetic documents and the corresponding segmentations. (1) Input synthetic documents. (2) Candidate bounding boxes obtained by parsing the PDF rendering commands. (3) Raw segmentation outputs. (4) Segmentations after post-processing. Segmentation label colors are: **figure**, **table**, **section heading**, **caption**, **list** and **paragraph**.

The Business of Software; Phillip G. Armour

### The Operational Executive Sponsor

The sometimes-missing link in software process change.


In his seminal work, *Managing the Software Process*, Walter J. Parekh outlined what he called the "Operational Executive Sponsor" role. This role is the link between the business and the software process. It is the person who is responsible for ensuring that the software process is aligned with the business goals. This role is often overlooked, but it is critical for the success of any software project.

... (text continues) ...



Block walls for extended periods of time, longer than had...

... (text continues) ...




The horizontal add XP operators are given in the table below:

Operator	Symbol	Meaning
Horizontal Add (Signed)	+	...
Horizontal Add (Unsigned)	+	...
Horizontal Subtract (Signed)	-	...

Multiple (M) Accumulate

... (text continues) ...




7.0 CURTIBA CASE STUDY

INTRODUCTION

Curtiba, Brazil is a medium size city in a developing country that through visionary planning and inspired leadership has adopted high investment yet low cost strategies to cope with rapid growth. The curitiba, multifunctional development of a world class urban transit network in close coordination with physical land use planning and Curitiba's rapid FID has set a new world class example, the success of Curitiba are important and insightful.

... (text continues) ...




B. Museums and Galleries

C. Architecture

D. Housing

A. Transport

... (text continues) ...



WEST CHESTER TOWNE CENTRE

The Square at Union Centre

... (text continues) ...




TABLE 12: Comparative performance by region, 1990-1999

A. Growth rates by region (% per year)

Region	GDP	Population	GDP per capita
Northwestern Europe	2.92	0.43	5.29
Southern Europe	1.40	0.89	1.55
Central & eastern Europe	1.48	0.91	2.35
Total Europe	1.60	0.89	2.32

B. Regional shares (% of total)


Region	1990	1999
Northwestern Europe	26.4	23.3
Southern Europe	23.4	23.4
Central & eastern Europe	49.1	53.3

C. Relative per capita income by region (% of European average)

Region	1990	1999
Northwestern Europe	132.9	130.3
Southern Europe	51.9	51.9
Central & eastern Europe	51.4	51.4

D. Growth rates for alternative groupings (% per year)

Grouping	GDP	Population	GDP per capita
OECD	2.06	0.43	4.75
BRIC	3.44	0.53	6.49
Eastern bloc countries	2.20	1.00	2.20
Developing countries	3.02	0.75	4.02



TIGERS FRONT OFFICE

... (text continues) ...




Figure 5: Real documents and the corresponding segmentations. Segmentation label colors are: **figure**, **table**, **section heading**, **caption**, **list** and **paragraph**.