

Final exam in 1DL301 Database Design I

Department of Information Technology, Uppsala University

January 10, 2020, 14.00 – 19.00

This is a multiple-choice exam with two types of questions:

- If a **question is marked with ♣** you must **select all correct choices**. If you don't select all correct choices or you include any incorrect choice, your answer will be marked as incorrect. There is always at least one correct choice in each such question.
- For all **other questions** you must **select only one choice** even if there are several correct choices. Your answer will be marked as correct if you select any of the correct choices. If you select an incorrect choice or select more than one choice, your answer will be marked as incorrect.

The questions are divided into four sections:

Section	Questions
Data modeling	10
SQL	10
FDs, CKs, NFs and normalization	10
Other	6

To pass the exam you need:

- to answer correctly at least **50% of the questions in each section, and**
- **at least 22 correct answers in total.**

Failing to meet either of these criteria means failing the exam (i.e. grade U).

If you fulfill both criteria, your grade will be determined by the number of correct answers:

Correct answers	Grade
22 – 26	3
27 – 31	4
32 – 36	5

Your answers must be given on the answer sheet which will be handed in. Don't forget to fill out your exam code. To mark your answer **fill the answer box entirely using a dark colored pen** (black or blue). Do not use a pencil! The optical character recognition system will not recognize ticks, crosses, circles or any other additional notes.

If you make a mistake on the answer sheet, request a new one and make sure you hand in the correct sheet (if you hand in several answer sheets, your exam will not be graded).

You can keep the question sheets. Consider noting your answers on these first and filling out the answer sheet just before handing it in.

Allowed aids One A4 sheet with handwritten notes (both sides can be used) which must be handed in with your exam (remember to write your exam code in one of the corners). An English explanatory dictionary and/or a translation dictionary between English and your mother tongue.

1 Information about you

Question i (Not assessed.) What is your study program?

- ☐ A Bioinformatics
- ☐ B DV or IT
- ☐ C STS
- ☐ D X
- ☐ E Other

Question ii (Not assessed.) In which period were you registered on the course for the first time?


- ☐ A Fall 2019, period 2
- ☐ B Fall 2019, period 1
- ☐ C Earlier period

Question iii (Not assessed.) How many lectures have you attended?

- ☐ A None or very few
- ☐ B Around 25%
- ☐ C Around 50%
- ☐ D Around 75%
- ☐ E Almost all or all

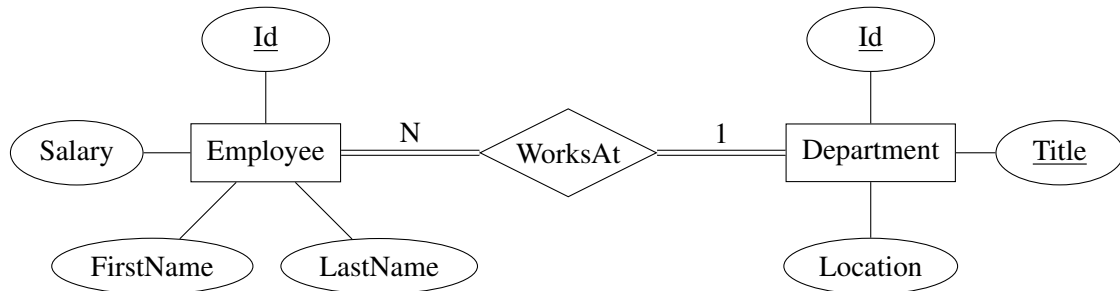
2 Data modeling

Question 1 What ER element is depicted in the following figure?



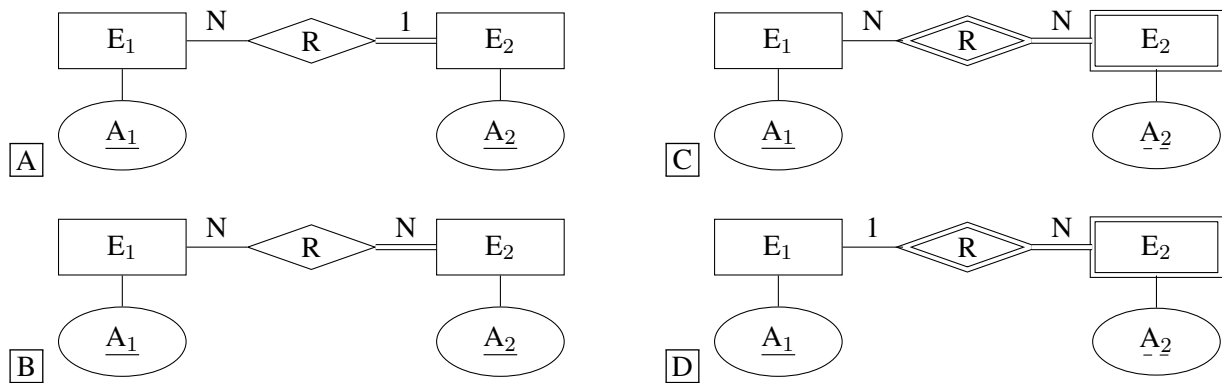
- ☐ A Multi-valued attribute
- ☐ B Weak attribute
- ☐ C Derived attribute
- ☐ D Composite attribute
- ☐ E Key attribute

Question 2 ♣ Based on the ER model depicted below and the semantics of ER models, which of the following statements are correct? (Select all correct choices.)




- ☐ A Each employee has a unique combination of the first and last name.
- ☐ B Each employee works at exactly one department.
- ☐ C Each department has at least one employee.
- ☐ D Each department has at least two employees.
- ☐ E Departments can have the same title as long the id is unique.
- ☐ F An employee can work at several departments.


Question 3 Which of the following ER models is not valid?







Question 4 Consider the following ER model:



Which of the following diagrams using the min-max notation depicts the same model?







Question 5 You have been asked to implement a simple time-tracking application, i.e. software that allows employees to record time spent on projects. At the end of each working day, all employees will record how many hours they have spent on each project they have been working on during that day, including a short summary of what was done. Which of the following ER diagrams shows a correct way to store the time-tracking data? (Only the relevant portion of each diagram is shown.)




Question 6 Which of the ER models corresponds to the following relational model?

$A(A_1, A_2)$

$B(B_1, B_2, A_1)$ with $A_1 \xrightarrow{FK} A(A_1)$




Question 7 Convert the following ER model to a relational model (foreign key constraints are not shown).







- ☐ A $E(\underline{E_1}, E_2), F(\underline{F_1}, F_2), R(\underline{E_1}, \underline{F_1})$
- ☐ B $F(\underline{E_1}, \underline{F_1}, E_2, F_2)$
- ☐ C $E(\underline{E_1}, E_2), F(\underline{E_1}, \underline{F_1}, F_2)$
- ☐ D $E(\underline{E_1}, E_2), F(E_1, \underline{F_1}, F_2)$
- ☐ E None of the other answers

Question 8 Convert the following ER model to a relational model.




- ☐ A $E(\underline{E_1}, E_2)$
 $F(\underline{F_1}, F_2)$
 $R(\underline{E_1}, \underline{F_1}, R_1)$ with $E_1 \xrightarrow{FK} E(E_1)$ and $F_1 \xrightarrow{FK} F(F_1)$
- ☐ B $E(\underline{E_1}, E_2)$
 $F(\underline{F_1}, F_2)$
 $R(\underline{E_1}, \underline{F_1}, R_1)$ with $E_1 \xrightarrow{FK} E(E_1)$ and $F_1 \xrightarrow{FK} F(F_1)$
- ☐ C $E(\underline{E_1}, E_2)$
 $F(\underline{F_1}, F_2)$
 $R(R_1)$
- ☐ D None of the other answers
- ☐ E $E(\underline{E_1}, E_2, F_1, R_1)$ with $F_1 \xrightarrow{FK} F(F_1)$
 $F(\underline{F_1}, F_2, E_1, R_1)$ with $E_1 \xrightarrow{FK} E(E_1)$

Question 9 A company has two types of projects: X and Y. Each project is either of type X or Y, but not both at the same time. Since many (but not all) of the attributes we wish to store for both project types are the same (such as the id, title, department, project manager, start date, etc.), we want to use a superclass to store these attributes. Which of the following EER models corresponds to this specification? (The attributes are not shown.)



Question 10 Convert the following ER diagram to a relational model.



A $P(\underline{P_1}, E_1, F_1)$

B $P(\underline{P_1})$
 $E(\underline{E_1})$
 $F(\underline{F_1})$

C $P(\underline{P_1})$
 $E(\underline{P_1}, \underline{E_1})$ with $P_1 \xrightarrow{FK} P(\underline{P_1})$
 $F(\underline{P_1}, \underline{F_1})$ with $P_1 \xrightarrow{FK} P(\underline{P_1})$

D $P(\underline{P_1})$
 $E(\underline{P_1}, E_1)$ with $P_1 \xrightarrow{FK} P(\underline{P_1})$
 $F(\underline{P_1}, F_1)$ with $P_1 \xrightarrow{FK} P(\underline{P_1})$

E $P(\underline{P_1})$
 $E(\underline{P_1}, \underline{E_1})$ with $P_1 \xrightarrow{FK} P(\underline{P_1})$
 $F(\underline{P_1}, \underline{F_1})$ with $P_1 \xrightarrow{FK} P(\underline{P_1})$

3 SQL

For questions 11 – 16 consider the following database consisting of tables T1 and T2:

T1	
C1	C2
A	A
A	B
B	C
D	D

T2	
C1	C2
A	D
A	A
B	A
C	NULL

Question 11 How many rows are in the result of the following SQL query?

```
SELECT * FROM T2 WHERE C2 = NULL
```

- ☐ A 1
- ☐ B 0
- ☐ C 3
- ☐ D 2
- ☐ E 4

Question 12 What is the result of the following SQL query? (Not showing the header of the result.)

```
SELECT DISTINCT T1.C1, T1.C2
FROM T1 JOIN T2 ON T1.C2 = T2.C1
WHERE T2.C2 IS NOT NULL;
```

☐ A

A	A
A	A
A	B

☐ B Empty table

☐ C

A	A
A	B

☐ D

A	A
A	A
A	B
B	C

☐ E

A	A
A	B
B	C

Question 13 What is the result of the following SQL query? (Not showing the header of the result.)

```
SELECT T1.C1, COUNT(T2.C2)
FROM T1, T2
WHERE T1.C2 = T2.C1
GROUP BY T1.C1
```

☐ A

A	2
A	1
B	0

☐ D

A	3
B	1

☐ B

A	3
B	0

☐ E

A	3
---	---

☐ C SQL is invalid

☐ F

A	2
B	0

Question 14 What is the result of the following SQL query? (Not showing the header of the result.)

```
SELECT T1.C1
FROM T1, T2
WHERE T1.C2 = T2.C1 AND COUNT(T2.C2) >= 1
GROUP BY T1.C1
```

☐ A SQL is invalid

☐ C

A
A

☐ B

A
B

☐ D Empty table

☐ E

A

Question 15 What is the result of the following SQL query? (Not showing the header of the result.)

```
SELECT T1.C1, T1.C2
FROM T2 RIGHT JOIN T1 ON T1.C2 = T2.C1
WHERE T2.C2 = "D" OR T2.C2 IS NULL
```

☐ A

A	A
B	C
D	D

☐ D Empty table

☐ B

A	A
B	C

☐ E

A	A
B	C
D	NULL

☐ C

A	A
D	D

☐ F SQL is invalid

Question 16 Some database management systems do not support RIGHT JOIN. For example, if you try to run the query from the previous question in SQLite, you will get the following error:

Error: RIGHT and FULL OUTER JOINS are not currently supported

What query can you run instead and get the correct result?

- ☐ A It is not possible to use a single query to get the same result.
- ☐ B

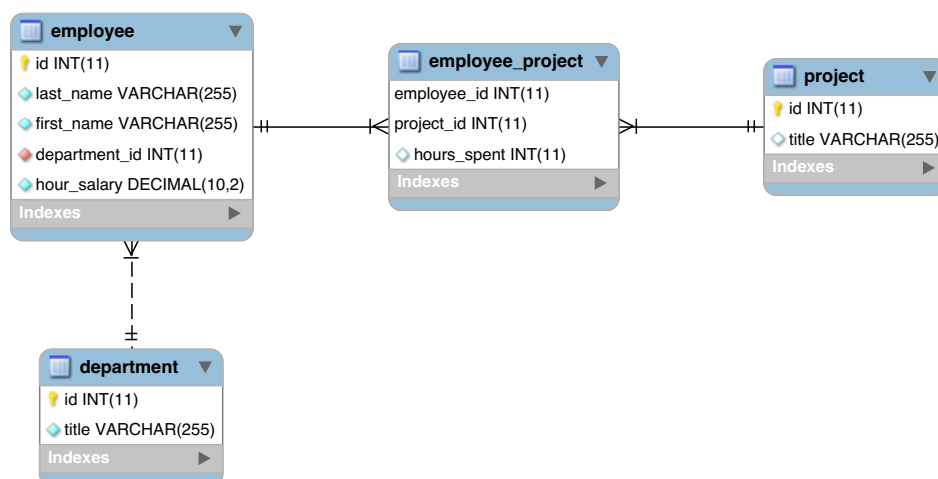
```
SELECT T1.C1, T1.C2
FROM T1 INNER JOIN T2 ON T1.C2 = T2.C1
WHERE T2.C2 = "D" OR T2.C2 IS NULL
```
- ☐ C

```
SELECT T1.C1, T1.C2
FROM T1 LEFT JOIN T2 ON T1.C2 = T2.C1
WHERE T2.C2 = "D" OR T2.C2 IS NULL
```
- ☐ D

```
SELECT T1.C1, T1.C2
FROM T1, T2
WHERE T1.C2=T2.C1 AND (T2.C2 = "D" OR T2.C2 IS NULL)
```
- ☐ E

```
SELECT T1.C1, T1.C2
FROM T1 JOIN T2 ON T1.C2 = T2.C1
WHERE T2.C2 = "D" OR T2.C2 IS NULL
UNION
SELECT C1, C2
FROM T1
WHERE T1.C2 NOT IN (SELECT T2.C1 FROM T2)
```

For questions 17 – 20 consider the following database model:



Question 17 ♣ Finish the SQL statement to create the employee_project table:

```
CREATE TABLE employee_project (
    employee_id int(11) NOT NULL,
    project_id int(11) NOT NULL,
    hours_spent int(11),
    ①
)
```

In the following list select **all** relevant primary and foreign key constraints that together (separated by commas) will replace ①.

- ☐ A FOREIGN KEY (project_id) REFERENCES project(id)
- ☐ B PRIMARY KEY (employee_id, project_id, hours_spent)
- ☐ C FOREIGN KEY employee(id) REFERENCES employee_id
- ☐ D PRIMARY KEY (project_id)
- ☐ E FOREIGN KEY project(id) REFERENCES project_id
- ☐ F PRIMARY KEY (employee_id)
- ☐ G PRIMARY KEY (employee_id, project_id)
- ☐ H FOREIGN KEY (employee_id) REFERENCES employee(id)

Question 18 ♣ Finish the SQL query to return the id and the title of each department together with the number of its employees:

```
SELECT d.id, d.title, COUNT(e.id)
FROM department d, employee e
①
②
```

Select one choice for ① and one choice for ②.

- ☐ A ① No WHERE clause necessary
- ☐ B ① WHERE e.department_id=d.id
- ☐ C ① WHERE e.id=d.id
- ☐ D ① WHERE e.department_id=d.title
- ☐ E ② GROUP BY department
- ☐ F ② GROUP BY d.id, d.title
- ☐ G ② GROUP BY employee
- ☐ H ② GROUP BY e.id

Question 19 Which of the following SQL queries selects the last name and the hour salary of the worst paid employees in each department?

- ☐ A `SELECT e.last_name, MIN(e.hour_salary) FROM employee e GROUP BY e.department_id`
- ☐ B `SELECT e.last_name, e.hour_salary FROM employee e WHERE EXISTS (SELECT * FROM employee WHERE hour_salary < e.hour_salary AND department_id = e.department_id)`
- ☐ C `SELECT e.last_name, e.hour_salary FROM employee e WHERE e.hour_salary = (SELECT MIN(hour_salary) FROM employee WHERE department_id=e.department_id)`
- ☐ D `SELECT last_name, hour_salary FROM employee GROUP BY department_id HAVING hour_salary=MIN(hour_salary)`
- ☐ E `SELECT e.last_name, e.hour_salary FROM employee e WHERE NOT EXISTS (SELECT * FROM employee WHERE hour_salary <= e.hour_salary)`

Question 20 Finish the query to select the id of employees working on either project 14 or project 15 (or both):

```
SELECT DISTINCT employee_id
FROM employee_project
WHERE ①
```

- ☐ A `project_id = 14 OR project_id=15`
- ☐ B `project_id = 14 AND project_id=15`
- ☐ C `project_id = 14 UNION project_id=15`
- ☐ D None of the other answers
- ☐ E `project_id IN (14, 15)`

4 FDs, CKs, NFs and normalization

For questions 21 – 25 consider a relation R(A, B, C, D, E) in 1NF with the following dependencies:

- $\{A, B\} \rightarrow \{C\}$
- $\{B\} \rightarrow \{E\}$
- $\{E\} \rightarrow \{D\}$

Question 21 ♣ Select all elements of $\{B, C\}^+$ (i.e. the set of all attributes which can be determined by the set $\{B, C\}$)?

- ☐ A A
- ☐ B B
- ☐ C C
- ☐ D D
- ☐ E E

Question 22 ♣ Select all candidate keys of R.

- ☐ A {B}
- ☐ B {A, B, E}
- ☐ C {E}
- ☐ D {A, B}
- ☐ E {A}

Question 23 ♣ Select all non-prime attributes.

- ☐ A A
- ☐ B B
- ☐ C C
- ☐ D D
- ☐ E E

Question 24 In which normal form (NF) is R?

- ☐ A BCNF but not 3NF
- ☐ B 3NF but not BCNF
- ☐ C 1NF but not 2NF
- ☐ D BCNF
- ☐ E 2NF but not 3NF

Question 25 Which of the following options shows the result of the normalization of R with all tables in BCNF (not showing the primary and foreign key constraints)?

- ☐ A $R_1(A, B, C), R_2(B, D, E)$
- ☐ B $R(A, B, C, D, E)$
- ☐ C $R_1(A, B, C, D, E), R_2(B, E), R_3(D, E)$
- ☐ D $R_1(A, C), R_2(B, D, E)$
- ☐ E $R_1(A, B, C), R_2(B, E), R_3(D, E)$

For questions 26 – 28 consider the following relation in 1NF:

`Student(StudentId, FirstName, LastName, DateOfBirth, Program, Campus)`

with the following dependencies:

- $\{StudentId\} \rightarrow \{FirstName, LastName, DateOfBirth, Program, Campus\}$
- $\{Program\} \rightarrow \{Campus\}$

Question 26 ♣ Select all choices that are superkeys of the Student relation.

- ☐ A {FirstName, LastName, DateOfBirth}
- ☐ B {StudentId, Program}
- ☐ C {StudentId}
- ☐ D {FirstName, LastName, DateOfBirth, Program, Campus}
- ☐ E {Program}

Question 27 ♣ Select all choices that are candidate keys of the Student relation.

- ☐ A {FirstName, LastName, DateOfBirth, Program, Campus}
- ☐ B {FirstName, LastName, DateOfBirth}
- ☐ C {StudentId}
- ☐ D {StudentId, Program}
- ☐ E {Program}

Question 28 In which NF is the Student relation?

- ☐ A 3NF but not in 2NF
- ☐ B 1NF but not in 2NF
- ☐ C BCNF
- ☐ D 2NF but not in 3NF
- ☐ E 3NF but not in BCNF

Question 29 Which of the following options shows the result of the normalization of the original table with all tables in BCNF (not showing the foreign key constraints)?

- ☐ A Student(StudentId, FirstName, LastName, DateOfBirth)
Program(StudentId, Program, Campus)
- ☐ B Student(StudentId, FirstName, LastName, DateOfBirth, Program, Campus)
Program(Program, Campus)
- ☐ C Student(StudentId, FirstName, LastName, DateOfBirth, Program)
Program(Program, Campus)
- ☐ D Student(StudentId, FirstName, LastName, DateOfBirth, Program)
Program(Program, Campus)
- ☐ E Student(StudentId, FirstName, LastName, DateOfBirth)
Program(StudentId, Program, Campus)

Question 30 The following table shows the current state of a relation:

C_1	C_2	C_3	C_4
1	1	XX	157
2	1	XX	178
3	1	XX	192
4	2	XY	183
5	2	XY	166

Is the following dependency?

$$\{C_2\} \rightarrow \{C_3\}$$

- ☐ A Yes, since if C_2 is 1, then C_3 is XX, and if C_2 is 2, then C_3 is XY.
- ☐ B A dependency cannot be confirmed by just checking one state of the relation.
- ☐ C No, values of C_2 are not unique.
- ☐ D Yes, since all rows are distinct.
- ☐ E No, values of C_3 are not unique.

5 Other

Question 31 When transferring money from one bank account to another we need to run two SQL statements, one that reduces the balance of the first account and one that increases the balance of the second account. Which of the ACID properties guarantees that either both SQL statements get executed (and the money is transferred) or none does (not transferring the money at all)?

- ☐ A Durability
- ☐ B Atomicity
- ☐ C Isolation
- ☐ D Consistency
- ☐ E Integrity

Question 32 Minimal superkey is better known as what?

- ☐ A Primary key
- ☐ B Unique key
- ☐ C Candidate key
- ☐ D Foreign key
- ☐ E None of the other answers

Question 33 Consider a relation $R(A_1, A_2, \dots, A_n)$, and a subset of its attributes, $K \subseteq \{A_1, A_2, \dots, A_n\}$. If for any legal instance r and any two distinct tuples $t_1, t_2 \in r, t_1 \neq t_2$ it holds that $t_1[K] \neq t_2[K]$ (where $t[K]$ denotes the projection of t onto the attributes in K , i.e. the values in t but only for the attributes in K), then K is:

- ☐ A Candidate key
- ☐ B Determinant
- ☐ C Superkey
- ☐ D Foreign key
- ☐ E Primary key

Question 34 Alice creates the table X. Then, the following sequence of statements is executed, in the given order (the name of the user executing the statement is indicated at the beginning of each statement):

Alice: GRANT ALL ON X TO Bob WITH GRANT OPTION
 Bob: GRANT SELECT ON X TO Jim WITH GRANT OPTION
 Jim: GRANT SELECT ON X TO Tim
 Alice: REVOKE ALL ON X FROM Bob CASCADE

Which privileges (concerning the table X) does Tim have?

- ☐ A SELECT with the grant option
- ☐ B None
- ☐ C All privileges, but without the grant option
- ☐ D SELECT without the grant option
- ☐ E All privileges with the grant option

Question 35 Consider a relation $R(A, B, C, D, E)$ containing 10^7 records. A is the primary key, and B and C each contain 10^5 distinct values. The following prepared SQL is executed very often:

SELECT A, B FROM R WHERE C=?

Which of the following indexes will make the query executions faster?

- ☐ A Index on R(C, B)
- ☐ B Index on R(B, C)
- ☐ C Index on R(C)
- ☐ D Index on R(A, B, C)
- ☐ E Index on R(A, B)

Question 36 Consider the following relation R:

C1	C2
1	110
2	140
3	120
4	130


We execute the following statements:

```
CREATE TABLE X AS SELECT AVG(C2) FROM R;  
CREATE VIEW Y AS SELECT AVG(C2) FROM R;  
UPDATE R SET C2=C2*2;
```

What are the results of the following queries (not showing the headers)?

- ① SELECT * FROM X
- ② SELECT * FROM Y

- ☐ A ① 250 ② 125
- ☐ B ① 125 ② 125
- ☐ C ① 250 ② 250
- ☐ D CREATE VIEW is not a valid SQL statement
- ☐ E ① 125 ② 250

**Final exam in 1DL301 Database Design I – Answer sheet**

Department of Information Technology, Uppsala University

January 10, 2020, 14.00 – 19.00

Your exam code

.....

Encode the number in your exam code in the table on the right side. E.g., if your exam code is AB-2097-CDE, fill in the box no. 2 in the first column, the box no. 0 in the second, the box no. 9 in the third, and the box no. 7 in the last column. Fill in the boxes entirely!

<input type="text"/>	0	<input type="text"/>	0	<input type="text"/>	0	<input type="text"/>	0
<input type="text"/>	1	<input type="text"/>	1	<input type="text"/>	1	<input type="text"/>	1
<input type="text"/>	2	<input type="text"/>	2	<input type="text"/>	2	<input type="text"/>	2
<input type="text"/>	3	<input type="text"/>	3	<input type="text"/>	3	<input type="text"/>	3
<input type="text"/>	4	<input type="text"/>	4	<input type="text"/>	4	<input type="text"/>	4
<input type="text"/>	5	<input type="text"/>	5	<input type="text"/>	5	<input type="text"/>	5
<input type="text"/>	6	<input type="text"/>	6	<input type="text"/>	6	<input type="text"/>	6
<input type="text"/>	7	<input type="text"/>	7	<input type="text"/>	7	<input type="text"/>	7
<input type="text"/>	8	<input type="text"/>	8	<input type="text"/>	8	<input type="text"/>	8
<input type="text"/>	9	<input type="text"/>	9	<input type="text"/>	9	<input type="text"/>	9

Use a **dark colored pen** (blue or black). Do not use a pencil! To mark your answer **fill in the box entirely** (■)! The OCR software will not recognize ticks, crosses, circles, etc. **Do not make any additional notes** on this sheet!

If you make an error, ask for a new answer sheet!**i** A: B: C: D: E: **ii** A: B: C: **iii** A: B: C: D: E: **1** A: B: C: D: E: **2** A: B: C: D: E: F: **3** A: B: C: D: **4** A: B: C: D: **5** A: B: C: D: **6** A: B: C: D: **7** A: B: C: D: E: **8** A: B: C: D: E: **9** A: B: C: D: **10** A: B: C: D: E: **11** A: B: C: D: E: **12** A: B: C: D: E: **13** A: B: C: D: E: F: **14** A: B: C: D: E: **15** A: B: C: D: E: F: **16** A: B: C: D: E: **17** A: B: C: D: E: F: G: H: **18** A: B: C: D: E: F:
G: H: **19** A: B: C: D: E: **20** A: B: C: D: E: **21** A: B: C: D: E: **22** A: B: C: D: E: **23** A: B: C: D: E: **24** A: B: C: D: E: **25** A: B: C: D: E: **26** A: B: C: D: E: **27** A: B: C: D: E: **28** A: B: C: D: E: **29** A: B: C: D: E: **30** A: B: C: D: E: **31** A: B: C: D: E: **32** A: B: C: D: E: **33** A: B: C: D: E: **34** A: B: C: D: E: **35** A: B: C: D: E: **36** A: B: C: D: E: