INFS1200 Diagram Notation

Table of contents

What was sampled . . . 2
The palette . . . 3
Relation schemas written as text . . . 4
Relation instances and result tables . . . 4
Highlighting inside a table . . . 5
Foreign keys drawn as arrows . . . 5
Constraint-violation figures (/) . . . 6
ER diagrams: do they match the Chen conventions er-diagram-notation states? . . . 6
Instance-level ("ball and line") relationship diagrams . . . 8
ER-to-relational mapping figures . . . 8
Progressive schema build-up . . . 9
Relational algebra notation . . . 9
Set-operation Venn diagrams . . . 10
SQL presentation . . . 10
Division and double-negation figures . . . 11
Functional dependency notation . . . 12
Decomposition and normalisation figures . . . 13
Mathematical and logical symbols in use . . . 14
Transactions . . . 14
Figure sources and textbook lineage . . . 14
Inconsistencies . . . 14
Not established by the source . . . 16

Reference legend for **every diagram family in INFS1200 except ER/EER symbols and ER-to-relational mapping**, which are already covered by er-diagram-notation[1] and relational-mapping-notation[2]. This note covers relation schemas and instances drawn as tables, foreign-key figures, constraint-violation figures, relational-algebra notation, SQL presentation,

[1] courses/infs1200/er-diagram-notation.pdf
[2] courses/infs1200/relational-mapping-notation.pdf

functional-dependency notation, and normalisation/decomposition figures — plus a section confirming whether the ER diagrams really are drawn to the Chen conventions the existing note asserts.

Everything below was inferred from the course's own PDFs. Where a convention is stated by the course rather than merely practised, that is called out.

What was sampled

The course ships **two written guides**, both one page, and both already covered by the sibling notes: the *EER & ER Notation Guide* (naming only — entity/relationship names capitalised and space-free, attributes UpperCamelCase) and the *Relational Mapping Guide* (`Relation [attr, ...]` schema form, table names UpperCamelCase, attributes lowerCamelCase, `X.fk references Y.pk`). **Neither guide says anything about symbol shape, colour, line style, table presentation, algebra notation, FD notation or figure layout.** Every convention in this note beyond naming was inferred from the drawings.

Deck shorthand used in citations below:

Cite	File	Pages
L0	`infs1200-lecture-0.pdf` (Module 0: Introduction / DBMS)	64
L1	`infs1200-lecture-1.pdf` (Module 1: ER Model)	60
L2	`infs1200-lecture-2.pdf` (Module 2: Relational Model + Mapping)	94
L3	`infs1200-lecture-3.pdf` (Module 3: SQL)	191
L4	`infs1200-lecture-4.pdf` (Module 4: Design Theory)	114

I read **261 of the 523 deck pages as rendered images**, plus both guide pages in full, plus a plain-text extraction of all 523 pages used to locate the pages carrying drawings. Specifically:

- L0: pp. 42–48, 59–63 (12 of 64)
- L1: pp. 28–34, 41–53 (20 of 60)
- L2: pp. 5–36, 43–58, 61–74, 76–93 (80 of 94)
- L3: pp. 5–9, 46–60, 71–82, 86–104, 106–118, 145–161 (81 of 191)
- L4: pp. 8–27, 52–70, 84–112 (68 of 114)

I did **not** read L0 pp. 1–41 and 49–58, L1 pp. 1–27 and 35–40, L3 pp. 10–45, 61–70, 83–85, 119–144, 162–191, or L4 pp. 1–7, 28–51, 71–83. The unread stretches are mostly prose, quiz slides and the generative-AI section; sampling targeted pages that carry drawings.

There is **no Module 5 (Database Security) deck in the source set at all** — only roadmap bullets naming it (L0 p63, L3 p2, L4 p3). See *Not established*.

The palette

Colours are used systematically enough to be treated as notation. Sampled directly from the rendered PDFs, so these are exact:

Hex	Name used below	Meaning
`#51247A`	**UQ purple**	Entity-type fill; table header fill; section-bar fill; derived/answer text
`#962A8B`	**magenta**	Relationship-diamond fill
`#E6E6E4` / `#E7E6E6`	**light grey**	Attribute-oval fill; alternating table body row; definition box; "this one is wrong" query box
`#E6E0EF`	**pale lavender**	Schema-reference box; SQL syntax-template box; quiz-stimulus box; side-note box
`#D1F2D3`	**pale green**	A query being demonstrated as correct
`#2EA836`	**green**	badge; "this is legal / this is the good design"
`#E62645`	**crimson**	badge; foreign-key arrows; "this violates the constraint" outlines
`#C00000` / `#CB0000`	**dark red**	Emphasis on a term being defined, in running prose

Consistency: the purple/magenta/light-grey trio holds for **every** ER diagram in L1 and L2. The green- /crimson- pairing holds across L2, L3 and L4 — it is the single most consistent visual convention in the course. The pale-green query box appears throughout L3 only (SQL is only taught there).

Relation schemas written as text

Follows the written guide, and is used identically in all four content decks.

- `Relation [attr1, attr2, ...]`, space before the bracket, primary key underlined. L2 pp. 15, 22, 30, 89–93; L3 pp. 49, 51, 56, 60, 87, 92–95, 109, 111, 112, 115; L4 pp. 56–59, 63, 64, 66, 68, 85, 88, 90, 93, 106, 107. **Throughout the semester.**
- A composite primary key gets **one continuous underline spanning all its attributes**, not one underline per attribute — `Enrollment [sid, cid, department, grade]` with a single rule under the first three. L2 pp. 30, 86, 90–93; L3 p49 (`StarsIn [movieID, starID, role]`); L4 p66. **Consistent.**
- Foreign keys are written on their own line directly under the referencing relation: `Employee.department references Department.id`. Composite FKs use brace sets on both sides: `Account.{bankCode, branchNo} references Branch.{bankCode, branchNo}`. L2 pp. 27, 30, 57, 86, 90–93. **Consistent.**
- Schema listings shown as *reference material* for a query or example sit in a **pale-lavender box with a thin black border**. L2 pp. 15, 27, 30; L3 pp. 49, 51, 54, 56, 57, 59, 60, 87, 92–96, 104, 109, 111, 112, 115, 151–161. **Throughout Module 3.**

Relation instances and result tables

The single most-repeated figure in the course. Identical in L0, L2, L3 and L4.

- **Header row**: solid UQ purple `#51247A`, white bold text, one attribute name per cell.
- **Body rows**: alternate light grey `#E6E6E4` and white, starting with grey in row 1.
- **Borders**: thin black, on all cells.
- **Relation name**: bold black, *above* the table, centred. L0 p46; L2 pp. 8, 10–13, 15, 17, 18, 23, 25–29, 35, 36, 74; L3 pp. 88, 89, 93–95, 98, 103, 112, 118, 146–150, 159–161; L4 pp. 9, 10, 15, 24, 26, 27, 45, 53, 54, 60, 66. **Throughout the semester.**
- Alignment is left in every table observed; numbers are not right-aligned.
- Key attributes are **underlined in the purple header** when the figure is about keys: L2 p74 (`PK1`, `PK2`), L4 p66 (`studentID`, `courseID`). Elsewhere headers are not underlined. Used **twice** — treat as available, not required.
- A **query result table** uses exactly the same styling as a base-relation table; the only difference is that it carries no relation-name caption, and its headers are the projected expressions themselves — `COUNT(*)`, `AVG (salary)` wrapping over two lines. L3 pp. 49, 73–76, 79, 93–95, 103, 112. **Consistent in Module 3.**
- **Empty cells** in a multiple-choice distractor are drawn as a grey cell with no content (L1 p34, L3 p118). Seen twice.

Highlighting inside a table

A small, stable vocabulary:

Device	Meaning	Citation
Thin red ellipse drawn over a cell or column	"look at this value" — sometimes the topic, sometimes the offence	L2 pp. 8, 10, 25, 26, 31, 33
Heavy black rectangle outline around a whole row	the row under discussion	L2 pp. 31, 33; L4 p27
Pale cyan/light-blue fill on rows or a rounded cyan outline on cells	the rows that demonstrate the problem	L4 p15; L3 pp. 159–161
Pale tint on a column, matched between two tables	pairs a foreign-key column with the primary-key column it references	L2 pp. 28, 29
Thick black diagonal strike-through across a row	this candidate has been ruled out	L3 pp. 159–161
Green beside a row	this candidate qualifies	L3 pp. 160, 161
Red text in a result-table cell	a **spurious tuple** produced by a lossy join	L4 p18

Consistency: the red ellipse and the black row-outline recur across L2 and L4. The cyan devices are confined to L3 pp. 159–161 and L4 p15.

Foreign keys drawn as arrows

Two distinct arrow renderings, both meaning the same thing, both with the **arrowhead landing on the referenced attribute** (i.e. the arrow points FK → PK, from the child to the parent). This direction never varies.

- **Over tables**: thick crimson `#E62645` orthogonal (right-angled) elbow arrows drawn *above* the tables, from the FK column heading, up and across, down into the referenced PK column heading. L2 pp. 27, 29.
- **Inside a schema box**: thin arrows from the FK attribute to the referenced attribute. Drawn diagonally/curved in L3 pp. 49, 51, 56, 60, 93–95, 109, 111, 112, 115; drawn as black orthogonal elbows in L3 pp. 54, 57, 59, 73–76, 79, 80, 87, 92, 96, 97, 104. **Both idioms are used constantly in Module 3.**
- L2 p27 is the only slide carrying an **explicit legend line** for it: `Foreign Keys Attribute -- Referenced attribute`, in magenta text under the schema box. That legend is the authority for the direction.

- Textual constraint captions sit under the table as plain black prose (`Department.manager references Employee.id`), in addition to the arrows. L2 pp. 27, 28, 29.

Constraint-violation figures (/)

The course's standard way of showing "this is illegal, that is legal":

- The candidate tuple is drawn as a **detached one-row table** below the relation it would be inserted into.
- An **illegal** tuple gets a crimson `#E62645` outline and a **crimson circle with a white** ; a **legal** tuple gets a green outline and a **green circle with a white** . L2 pp. 21, 25, 26, 31, 33, 35, 36. **Consistent through the integrity-constraints section.**
- The same / badges are reused for anything that has a right and a wrong version: two candidate SQL queries (L3 pp. 104, 111, 113), two candidate decompositions (L4 pp. 15, 26, 27), two design approaches (L4 p109). In L3 the losing query's box is **light grey** and the winning one's is **pale green**; the badge sits at the box's top-left (L3 p104) or top-right (L3 pp. 111, 113). **Semester-wide for the badges; box-fill pairing is Module 3 only.**
- The five integrity-constraint *types* each get a fixed fill on a rounded "one-square-corner" banner, reused every time that type is discussed: domain = magenta-pink, key = gold/tan, entity integrity = cyan, referential integrity = grey, semantic integrity = UQ purple. L2 pp. 20, 21, 25, 26, 31, 32, 33. **Consistent within Module 2; never reused elsewhere.**

ER diagrams: do they match the Chen conventions er-diagram-notation[3] states?

Yes — with fills and colours the existing note does not record. The lectures carry their own two-slide **Notation Guide** (L1 pp. 50–51), which is the authority. Drawn form:

Symbol	Drawn as	Citation
Entity type	Solid **UQ purple** rectangle, **white** text, ALL CAPS name	L1 pp. 41–44, 46–49, 50, 51; L2 pp. 48, 50–53, 55, 56, 58, 66–72, 85, 88
Weak entity type	The same purple rectangle nested inside a second black rectangle (double border)	L1 pp. 41–44, 50; L2 pp. 48, 55, 56, 58, 88
Attribute	**Light grey** ellipse, thin black border, black text, joined by a thin plain line	L1 pp. 41–44, 47–50; L2 pp. 48, 50–53, 56, 58, 88
Key attribute	Same, name **solid-underlined**	L1 pp. 41–44, 48, 50; L2 pp. 48, 88

[3] courses/infs1200/er-diagram-notation.pdf

Symbol	Drawn as	Citation
Partial key attribute	Same, name **dash-underlined**	L1 pp. 41, 42, 44, 50; L2 pp. 48, 55, 56, 58, 88
Multivalued attribute	Light grey ellipse with a **double** black outline	L1 p50; L2 pp. 48, 50, 52
Derived attribute	Light grey ellipse with a **dashed** outline, joined by a **dashed** line	L1 p50; L2 pp. 48, 52
Relationship type	**Magenta** diamond, **white** text, ALL CAPS name	L1 pp. 43, 44, 48–51; L2 pp. 48, 50–53, 66–72, 85, 88
Identifying relationship	Magenta diamond nested inside a second diamond outline	L1 pp. 43, 44, 50; L2 pp. 48, 55, 58, 88
Cardinality ratio	Plain black `1` / `N` / `M`, placed on the edge near the entity end	L1 pp. 31–34, 41–43, 51; L2 pp. 48, 88
Total participation	**Double line** on that entity's edge into the diamond	L1 pp. 42, 43, 51; L2 pp. 48, 62, 88
Subclass	drawn as an **open arc** on the connecting line, with the arc's opening facing the **superclass**	L1 pp. 46–49, 51; L2 pp. 48, 85
Several subclasses	A small **light-grey circle** joining the subclass lines to the superclass; each subclass line carries its own arc	L1 pp. 48, 49, 51; L2 pp. 48, 85
Disjoint / overlapping	`d` / `o` inside that circle	L1 pp. 48, 49, 51; L2 pp. 48, 85
Total / partial specialisation	**Double** / single line from the superclass down to the circle	L1 p49 (legend + figure)
Role name (recursive relationship)	Plain lowercase black text beside the line	L2 p48 (`supervisor`, `supervisee`)

So the existing note's symbol list is correct; it just does not record that entities are filled purple, relationships filled magenta, and attributes filled grey — which is what makes a hand-drawn INFS1200 diagram look like an INFS1200 diagram. That is also why the Mermaid approximations described at the bottom of er-diagram-notation[4] are a genuine substitute rather than a reproduction: Mermaid cannot produce the double-outline shapes, the dashed ellipse, the underline styles, or the arc.

[4] courses/infs1200/er-diagram-notation.pdf

Layout: no fixed flow direction. Entity boxes sit on a rough horizontal spine with attributes fanned above and below (L1 pp. 41–44; L2 p88); specialisation hierarchies run top-down with the superclass on top (L1 pp. 47–49).

Instance-level ("ball and line") relationship diagrams

Used only to explain cardinality ratios, and the ellipse outline colour is **semantic**:

- Two (or three) vertical ellipses, each holding dots for entity instances, joined by thin black straight lines for relationship instances.
- Ellipse outline colour encodes the ratio, and the same colour is reused as an inline badge on the worked examples: **purple `#51247A` = 1:1**, **orange `#EB602B` = N:1**, **blue `#4085C6` = 1:N**, **green `#2EA836` = M:N**. Defined on L1 p30 and reused on pp. 31, 32, 33. **Consistent across four slides, and used nowhere else.**
- L1 p29 uses a different idiom for the same idea: purple person/item icons with **green arrows inside a solid green box** for what is allowed and **red arrows inside a dashed red box** for what is not. Single instance.

ER-to-relational mapping figures

- Slides are two-panel: **"ER Diagram"** (bold black heading) on the left/top, **"Mapping to Relation"** on the right/bottom. L2 pp. 50–53, 55, 56, 61, 62, 65–68, 71, 72, 76, 77, 83–85. **Consistent through the mapping section.**
- The **rule box** stating the step is a rounded rectangle with a solid fill and white text, and **its fill colour tracks the step group** shown on the roadmap slide (L2 p47): UQ purple for steps 1–2 (entities/weak entities), gold/tan for steps 3–5 (binary relationships), magenta for steps 6–7 (multivalued / N-ary), slate for the EER subclass step. L2 pp. 50, 55, 61, 65, 71, 76, 78, 83. **Consistent — the colour is load-bearing, not decorative.**
- The **"Mapping to Relation"** panel is a rounded rectangle with a purple border and white fill. What the current step *adds* is written in **red**; what was already there stays black. L2 pp. 50, 55, 61, 65–68, 71, 72, 76, 77,

  85. (One exception — see Inconsistencies.)

- **Explanatory side notes** sit in a pale-lavender rounded rectangle with a purple border. L2 pp. 51–53, 56, 62.
- Diagram annotations (e.g. "Total Participation") are black text in a purple-bordered white box, joined to the thing they label by a purple wedge pointer. L2 p62. Single instance.
- Lecturer commentary on a diagram is **red text with thin red arrows** — L1 p44, L2 pp. 59, 74. **Recurs; the standard callout style.**

Progressive schema build-up

The course builds a schema across consecutive slides and marks "what is new" — but **in two different ways**. See Inconsistencies; both are recorded here:

- **Grey-out idiom**: previously-established lines are **grey**, the newly-added lines are **black**, foreign-key lines are *italic*. L2 pp. 54, 57, 63, 69, 73.
- **Bold idiom**: everything stays black and the newly-added lines are **bold**. L2 pp. 89–93.

The "Final Schema" summary is enclosed in a **crimson rounded-rectangle outline**, laid out in two columns, with all foreign-key lines in italic. L2 p86. Single instance, but it is the canonical worked answer.

The same grey-out idiom marks progress through the normal forms: on the "Normalisation – Overview" slide all four normal forms are listed and every one except the current is greyed. L4 pp. 49, 50, 51, 55, 61, 67, 80. **Consistent.**

Relational algebra notation

The course teaches relational algebra as *motivation* for SQL, not as a calculus you write. Only some operators get symbols:

Operator	Symbol as written	Citation
Join / natural join	⋈ (bowtie), written infix: `R ⋈ R`	L3 pp. 90, 91 ("Denoted as ⋈"); L4 pp. 16, 17
Cartesian product	`X` — a plain capital letter X, *not* `×`: `R1 X R2`	L3 pp. 88, 89, 91 ("denoted as X")
Theta-join condition	`{=, ≠, <, >, ≤, ≥}`	L3 p99
Union / intersection / difference	`∪` / `∩` / `−`, named as "the relational algebra operations"	L3 pp. 106, 107, 116
Division	Written both as `A/B` (`R1/R2`) in prose and as `÷` in figures	L3 pp. 145–151
Relation cardinality	vertical bars around the relation name, e.g. `|R1| = m`, `|R1 X R2| = m * n`	L3 p88
Result of a decompose-then-join round trip	`R'` (prime)	L4 p18

There is **no ⋈ and no ×** anywhere in the sampled pages. Selection and projection are introduced purely as SQL clause names — "Projection (SELECT clause)" (L3 p48), "Selection

(WHERE clause)" (L3 p55) — never as algebraic operators. Do not write `_{cond}(R)` or `_{attrs}(R)` in an INFS1200 note; the course has no such notation.

Operator figures — a relational-algebra operation applied to concrete tables is drawn one of two ways:

- Operand tables at left, a **thick UQ-purple block arrow** between them and the result, with the expression (`R  R`) as a black label above the arrow. L3 p90; L4 pp. 16, 18. **Consistent.**
- Operand tables side by side with the **operator glyph set large between them** and the result table after an `=`. L3 pp. 88 (X), 148 (÷), 151 (÷).
- Relation names for schematic tables (`R`, `R` , `R` , `S`, `A`, `B1`, `B2`) are placed **to the left of the table** (L3 pp. 90, 149, 150; L4 pp. 14, 16,

  18) or **centred above it** (L3 pp. 103, 146–148; L4 p18 for `R` and `R'`). Both are common.

Set-operation Venn diagrams

Two incompatible idioms, both in Module 3:

- **Black outline, diagonal hatching** for the selected region, in a bordered 3-row legend table (Union / Intersection / Difference). L3 p106.
- **Colour-filled circles** (yellow intersection, magenta left, orange right) for the join types, with lowercase function-style captions `inner_join(x, y)`, `left_join(x, y)`. L3 p102 — visibly an imported graphic (dplyr-style), not the course's own drawing.

SQL presentation

This is where the course is most systematic, and it states its own legend.

Syntax templates (how a statement is *specified*) sit in a **pale-lavender** box, and L3 p72 gives the full legend explicitly in its own box:

```
Notation
  KEYWORD          reserved word, ALL CAPS, coloured red
  <argument>       a placeholder you replace
  [optional]       may be omitted
  {multiple}       may be repeated
  …|..choice..|…   pick one
```

Abbreviated versions of the same legend appear on L3 pp. 48 and 55. Keywords in these templates are **red and bold**. **Consistent for every syntax box in Module 3** (L3 pp. 48, 55, 72, 91, 108, 110, 114).

Worked queries (an actual query being run) sit in a **pale-green `#D1F2D3`** box with a thin dark border. Keywords are plain black, not red, and the query is laid out in two columns: the clause keyword in a left column, its arguments in a right column, one clause per line. L3 pp. 49, 51, 56, 57, 59, 60, 73–76, 79, 80, 87, 92–95, 97, 104, 109, 111–113, 115. **Consistent.**

Quiz-stimulus queries (a query you are asked *about*, not one being taught) sit in a **pale-lavender** box like the templates — L3 pp. 77, 82, 89, 96, 98, 100, 118. So box fill separates "here is how you write it" / "here is one that works" / "here is one to reason about".

Other SQL conventions:

- **Casing**: keywords ALL CAPS; table names UpperCamelCase; attribute names lowerCamelCase — matching the Relational Mapping Guide. L3 throughout.
- **Aliases**: a short uppercase letter or initialism (`D`, `E`, `S`, `MS`, `M`, `F1`, `F2`, `S1`, `S2`, `M1`, `M2`, `A`, `B`, `X`), introduced either bare (`Employee E`) or with `AS` (`Employee AS E`). L3 p87 explicitly says both are acceptable; both appear throughout.
- **Terminating semicolon** is present on some queries and absent on others, with no discernible rule (present L3 pp. 73, 79, 92, 104; absent L3 pp. 49, 56, 60, 75, 76, 93–95, 109, 115).
- **Selective red highlighting inside a black query** marks the one keyword the slide is about — `DISTINCT` (L3 p93), `<HAVING …>` (L3 p82), `INSERT INTO … SELECT` (L3 p77). Occasional, not systematic.
- Slides that are **live in-class exercises** carry a purple rounded-square icon of two figures with `</>` badges in the top-right. L3 pp. 49, 51, 54, 56, 57, 59, 60, 80, 89, 92, 94, 97, 104, 109, 112. **Consistent marker.**

Division and double-negation figures

The most elaborate figure family in the course, and internally very disciplined (L3 pp. 152–161):

- Slide split into two panels by a **vertical purple rule**, headed **"Query"** (left, the SQL) and **"Conceptual"** (right, a set picture). **Consistent across all ten slides.**
- The set picture is two large overlapping circles, thick outline, unfilled, with plain black text captions above each ("Movie IDs from 1934", "Movie IDs acted in by X"). Regions of interest are filled with a deliberately **hand-drawn cyan scribble** rather than a flat fill.
- **Circle outline colour is tied to a query fragment**: the left circle is crimson, the right circle is green (L3 pp. 154–161), and **coloured vertical bars in the left margin of the SQL** mark which block of the query computes which circle — a crimson bar

beside the crimson circle's subquery, a green bar beside the green circle's subquery, a cyan bar beside the highlighted region.

- The **source tables at the bottom carry matching header colours**: crimson header on `StarsIn JOIN Movie` (the crimson circle), green header on `StarsIn (M2)` (the green circle), purple header on ordinary sub-tables. L3 pp. 159–161. **This colour-linking is the clearest example of colour carrying meaning anywhere in the course.**
- **Progressive refinement**: a natural-language description inside a **cyan rounded rectangle** stands in for a subquery not yet written, and is replaced by real SQL on the next slide. L3 pp. 154–158.
- Keywords in this section are **bold black**, not the plain black of the green query boxes elsewhere in L3.

Functional dependency notation

- An FD is written `X → Y` with a single-headed arrow, inline in prose and as a standalone caption. L4 pp. 22–24, 26, 27, 56–59, 62, 63, 66, 68, 69, 84, 85, 87–107. **Throughout Module 4.**
- **Composite sides are braced, single attributes are bare**: `{houseNum, street, postcode} → {state, value}` alongside `postcode → state` in the same FD set. L4 pp. 56–59, 88, 90, 93, 96, 102, 104, 106, 107. **Consistent.**
- **An FD set is written `F = { fd, fd, ... }`**, one FD per line, opening brace on the `F =` line. L4 pp. 56–59, 63, 66, 68, 69, 88, 89, 94, 96, 102, 104, 106.
- **Closure**: `F` for the closure of a set of FDs; `X` for the closure of an attribute set; `{A, B} = {A, B, C}` for a worked one. L4 pp. 34–37, 55, 57, 58, 61, 63, 67, 68, 85, 88, 90, 93, 96, 99, 100. **Consistent.**
- **Colour in FD statements**: the determinant (LHS) is written in **purple**, the dependent (RHS) in **magenta/plum** — `level → salary`, `B → A`, `postcode → state`. L4 pp. 22, 24, 26, 57, 58, 62, 63, 68, 69, 85, 87, 90, 93. **Consistent whenever an FD is called out as the subject of the slide**; FDs appearing incidentally inside an `F = {...}` set stay black.
- Derived results — closures, candidate keys, resulting relation schemas — are written in **purple**. L4 pp. 57, 66, 68, 85, 90, 93, 106, 107. **Consistent.**
- The formal tuple-comparison form is `if t1[X] = t2[X] → t1[Y] = t2[Y]`, with square brackets for tuple projection. L4 p23; also `t [FK] = t [PK]` at L2 p28.
- **Definition boxes** are light-grey rectangles with a thin black border and the defined term in **bold dark red**; a matching **white** box with a bold black `Example` heading sits directly under it. L4 pp. 14, 16, 52, 56–58, 62, 63, 68, 88, 105. **Consistent through Module 4.**

Decomposition and normalisation figures

BCNF decomposition tree — the most distinctive Module 4 figure, and drawn identically every time:

- A binary tree running top-down. Each node is a **white rectangle with a thin purple border** containing `Rn [A, B, C]`.
- Parent joins its two children with two thin plain purple lines forming an inverted V.
- The **violating FD is written as a plain black label in the crotch of the V** — `B→C`, `{A, B}→C`, `D→E`, `A→B`.
- **Left child = `Q1 [Q - Y]`, right child = `Q2 [X  Y]`.** This ordering never varies. L4 pp. 85, 90, 91 (all four multiple-choice variants), 93. **Consistent — four slides plus four distractors.**
- The result is stated in a **light-grey, black-bordered banner** spanning the slide: `Final answer: R2 [B, C], R3 [D, B], R4 [D, A]`. L4 pp. 85, 90, 93.

The `Q1 / Q2` mnemonic: two horizontally overlapping ellipses, black outline, no fill; the left labelled `Q1` containing "Others" and `X`, the right labelled `Q2` containing `X` and `Y`, with `X` in the overlap in its own small ellipse. Reproduced verbatim on L4 pp. 84, 85, 87, 90, 92, 93. **Consistent.**

Minimal-cover step table: columns headed `Step 1` / `Step 2` / `Step 3` (/ `Final`) with a purple header and white bold text; body cells in pale mauve-grey `#D0CCD7`; one FD per line; the FDs **changed at that step** picked out. L4 pp. 98, 99, 100, 102.

Minimal-cover flow diagram (the "Visualised" slides, L4 pp. 101, 103) — a Sankey-like trace, and the arrow style is meaningful:

- Horizontal light-grey bands, one per step, labelled at the far left (`Step 1: RHS simplification`, …, `Final Answer`, `Final with Union`).
- Each FD is a small white rectangle with a black border.
- A **thin black arrow with an arrowhead** = this FD was *transformed* at this step (split into several, or shortened).
- A **thin grey vertical line with no arrowhead** = this FD passed through unchanged.
- A box labelled `Removed` is the sink for FDs deleted at Step 3. **Both slides use it identically.**

Algorithm / pseudocode boxes: light-grey fill, thin black border, plain proportional text using `:=`, `{ }` blocks, `While ( ) { }`, `For each X → Y in G { }`, , , . The line being emphasised on that slide is turned **bold red**. L4 pp. 84, 87, 92, 95, 105. **Consistent.**

Decomposition arrows: a thick UQ-purple block arrow labelled `decomposition` in red (L4 p14) or `Decompose` / `Join` in black (L4 p18); or a purple bracket fanning from the parent table down to the two child tables (L4 p26, L3-style). Two idioms for one idea.

Mathematical and logical symbols in use

Recorded so a new note does not invent alternatives: (proper subset, L4 pp. 56, 58, 63, 68, 88), (L4 pp. 56, 58, 63, 99), (L4 pp. 84, 86, 88, 95, 105), (L4 pp. 95, 105), (L4 p96), (entailment, L4 p93 — single occurrence), `F` , `X` , `G'` .

Transactions

Only two slides in the whole sampled set (L2 pp. 43, 44), and **neither carries a diagram** — the transaction concept is prose plus an SQL example. There is no transaction-diagram notation in this course.

Figure sources and textbook lineage

- The course textbook is named on L0 p8: **Elmasri & Navathe, *Fundamentals of Database Systems*** (edition not stated on the slide; linked via UQ Library). No figure caption in the sampled pages names a source, but several figures are visibly lifted from that textbook and are drawn in **its** style, not the course's:
  - The instance-level `EMPLOYEE / WORKS_ON / PROJECT` figures with grey-filled ellipses, italic `e … e` , `r … r` , `d … d` labels and square dots for relationship instances — L1 pp. 31, 32, 33.
  - The three-schema architecture figure (green "External View" boxes, black box with yellow "Conceptual Schema" text, navy "Internal Schema" bar, gold DB cylinders, blue mapping captions) — L0 pp. 60, 61, 62. Also the red/green/navy DBMS block diagram at L0 pp. 42, 43, 48.
  - The EER superclass/subclass figure with white-filled, purple-outlined rectangles and ellipses — L1 p47. These use palettes and shape conventions found nowhere else in the course. Treat them as quotations, not as the house style.
- **L0 p47** shows a **MySQL Workbench** EER model screenshot — crow's-foot / IE-style boxes listing typed columns. It is a screenshot of the practical tooling, not a notation the course asks you to draw.
- **L3 p102**'s join Venn diagrams are an imported graphic with `inner_join(x, y)`-style captions.

Inconsistencies

Recorded, not resolved. Each of these needs an explicit decision before the first diagram that touches it.

1. **How `NULL` is spelled.** Uppercase `NULL` (L2 pp. 12, 13, 27, 28, 29), lowercase `null` (L2 p74; L4 p54; L3 p91 body text), and title-case `Null` (L3 p103). All three within the same course.
2. **String literal quoting in SQL.** Double quotes — `gender = "Male"` (L3 p56), `M.title LIKE "Gone with the Wind"` (L3 p95). Single quotes — `WHERE gender = 'female'` (L3 p60), `'Giants'` (L3 p58), `b = 'X'` (L3 p118). Both appear in worked, correct queries.
3. **Attribute-name case in table headers vs. schema text.** The schema line says `Employee [id, name, salary, department]` (lowercase) while the very same slide's instance table has headers `ID Name Salary Department` (capitalised) — L2 p15, and the same split on pp. 27–29, 35, 36. L4 p60 goes further and mixes them inside one schema: `Employee [ID, name, level, salary]`. Module 3 is consistently lowerCamel in both places (L3 pp. 51, 60), but L3 p49's schema box uses `Movie [MovieID, Title, Year]` while L3 p51's uses `Movie [movieID, title, year]` — the same relation, two slides apart.
4. **The same relation's key attribute is named two ways.** `Department.id` in the FK caption but `Number` as the table's first column header — L2 pp. 31, 33 versus pp. 27, 28.
5. **Which colour means "newly added" in a mapping box.** Red on L2 pp. 50, 55, 61, 65–68, 71, 72; **green** on L2 p62 (`mgrSSN, mgrStartDate`), where purple simultaneously means "relation name" and green means "attribute name".
6. **Two idioms for progressive schema build-up.** Grey-out-the-old (L2 pp. 54, 57, 63, 69, 73) versus bold-the-new (L2 pp. 89–93). Same deck, same worked example genre.
7. **Whether FK constraint lines are italic.** Italic on L2 pp. 56, 62, 67, 68, 71, 72, 86 and on every "Schema (in progress)" slide; upright on L2 pp. 50, 55, 61, 65, 66 and on L2 pp. 90–93.
8. **The arrow → is overloaded.** It is the FD arrow everywhere in Module 4, but L4 pp. 59 and 64 also use it for a *foreign key*: `Address.postcode → Postcodes.postcode`, `StaffAppointment.level → StaffIncome.level` — where the guide and all of Module 2 say `X.fk references Y.pk`. Two different relations written with one glyph.
9. **Weak-entity and identifying-relationship double lines are drawn in two colours.** Black in the Notation Guide (L1 p50) and in L1 pp. 43, 44 and L2 pp. 48, 55, 58, 88 — but **crimson** on L1 pp. 41, 42, where the `INSURES` diamond is also filled **blue** instead of magenta.
10. **rendering.** An open arc on the line (L1 pp. 46–49, 51; L2 pp. 48, 85) versus small arrowhead-like glyphs pointing from the `d` circle down at the subclasses (L2 p83). Same meaning, two glyphs.
11. **FD attribute case in Module 4.** `level → salary` (L4 pp. 22, 24, 26) versus `Level → Salary` (L4 p27) — consecutive slides, same FD.
12. **What is emphasised in the minimal-cover step table.** Bold **red** for the changed FDs on L4 pp. 98–100; bold **black** on L4 p102.
13. **Minor label drift in the same figure repeated.** `Removed` (L4 p101) vs `removed` (L4 p103); `Final W/ Union` (p101) vs `Final with Union` (p103); `Step 3: FD set simplification` (p101) vs `Step 3: FD removal` (p102).
14. **Venn-diagram idiom.** Hatched black-and-white (L3 p106) versus colour-filled (L3 p102)

— and the "Conceptual" circles change colour mid-run: purple right-hand circle on L3 pp. 152, 153, green on pp. 157–161.

15. **Relation-name label placement.** Bold above the table (most of L2/L4) vs to the left of the table (L2 pp. 31, 33; L3 pp. 90, 149, 150; L4 pp. 14, 16).
16. **ER shapes in Module 4 are unfilled.** L4 pp. 9 and 22 draw an ER fragment with a **white rectangle with a black border and bold black text** and **white ovals** — not the purple/grey fills the Notation Guide mandates.
17. **Result-table colouring for a derived relation.** L3 p88 gives the Cartesian-product result a **blue** header and blue-tinted rows to distinguish it from its purple-headed operands; every other result table in the course is purple. One-off.
18. **Terminating semicolons on SQL** — present on some worked queries, absent on others, with no visible rule (see SQL section for citations).

Not established by the source

Where a future INFS1200 diagram will have to invent something. Be candid about these rather than quietly picking a convention.

- **Everything about Module 5 (Database Security).** There is no security deck in the source set — only roadmap bullets naming "Database Security Issues / Access control mechanisms / Other Control Measures / SQL Injection" (L0 p63, L3 p2, L4 p3). So there is **no** established notation for access-control matrices, authorisation/`GRANT`–`REVOKE` graphs, role hierarchies, discretionary-vs-mandatory access diagrams, or SQL-injection attack flows. Anything drawn for database-access-control[5], sql-injection[6] or 2026-05-11-database-security[7] is invention, and should say so.
- **Selection and projection as algebra.** No , no anywhere in the sampled pages. If a note needs to write a projection algebraically there is no course-sanctioned form; use SQL or prose instead.
- **Rename in algebra.** never appears; renaming is taught only as SQL `AS` / qualification (L3 pp. 86, 87).
- **A symbol for division.** ÷ appears in figures and `/` in prose, but the course never states which to write in an answer.
- **Transaction diagrams.** No schedule diagrams, no precedence graphs, no ACID/state figures — the transaction concept is prose only (L2 pp. 43, 44).
- **Concurrency, recovery, indexing, storage.** Named as DBMS functions in L0 pp. 50–55 but explicitly deferred to later courses (L0 pp. 43, 48); no figures.
- **Any notation for `n`-ary relationship instance sets.** Cardinality for ternary relationships is discussed (L2 pp. 78–80) but the ball-and-line instance diagram is only ever drawn for binary relationships (L1 pp. 30–33).

[5] courses/infs1200/database-access-control.pdf
[6] courses/infs1200/sql-injection.pdf
[7] courses/infs1200/2026-05-11-database-security.pdf

- **Whether a subclass box may carry its own key attribute** in a drawn EER diagram; the drawn examples only ever give subclasses non-key attributes (L1 pp. 47, 48; L2 p48).
- **Line-crossing / routing style in ER diagrams.** All sampled diagrams are small enough to avoid crossings; there is no established jump, gap or corner convention.
- **Figure captions and numbering.** No figure in any sampled deck carries a number or a caption in the "Figure N:" sense; figures are titled only by the slide heading. There is therefore no cross-reference convention to copy.
- **Fonts and stroke weights** are PowerPoint defaults (Arial-like sans throughout); no monospace font is used for SQL or pseudocode anywhere, which is worth knowing before a note renders SQL in a code font and diverges from the source.
- **Dark-mode / print-safe variants.** Every figure assumes a white background; the purple/magenta/grey palette has no stated alternative.