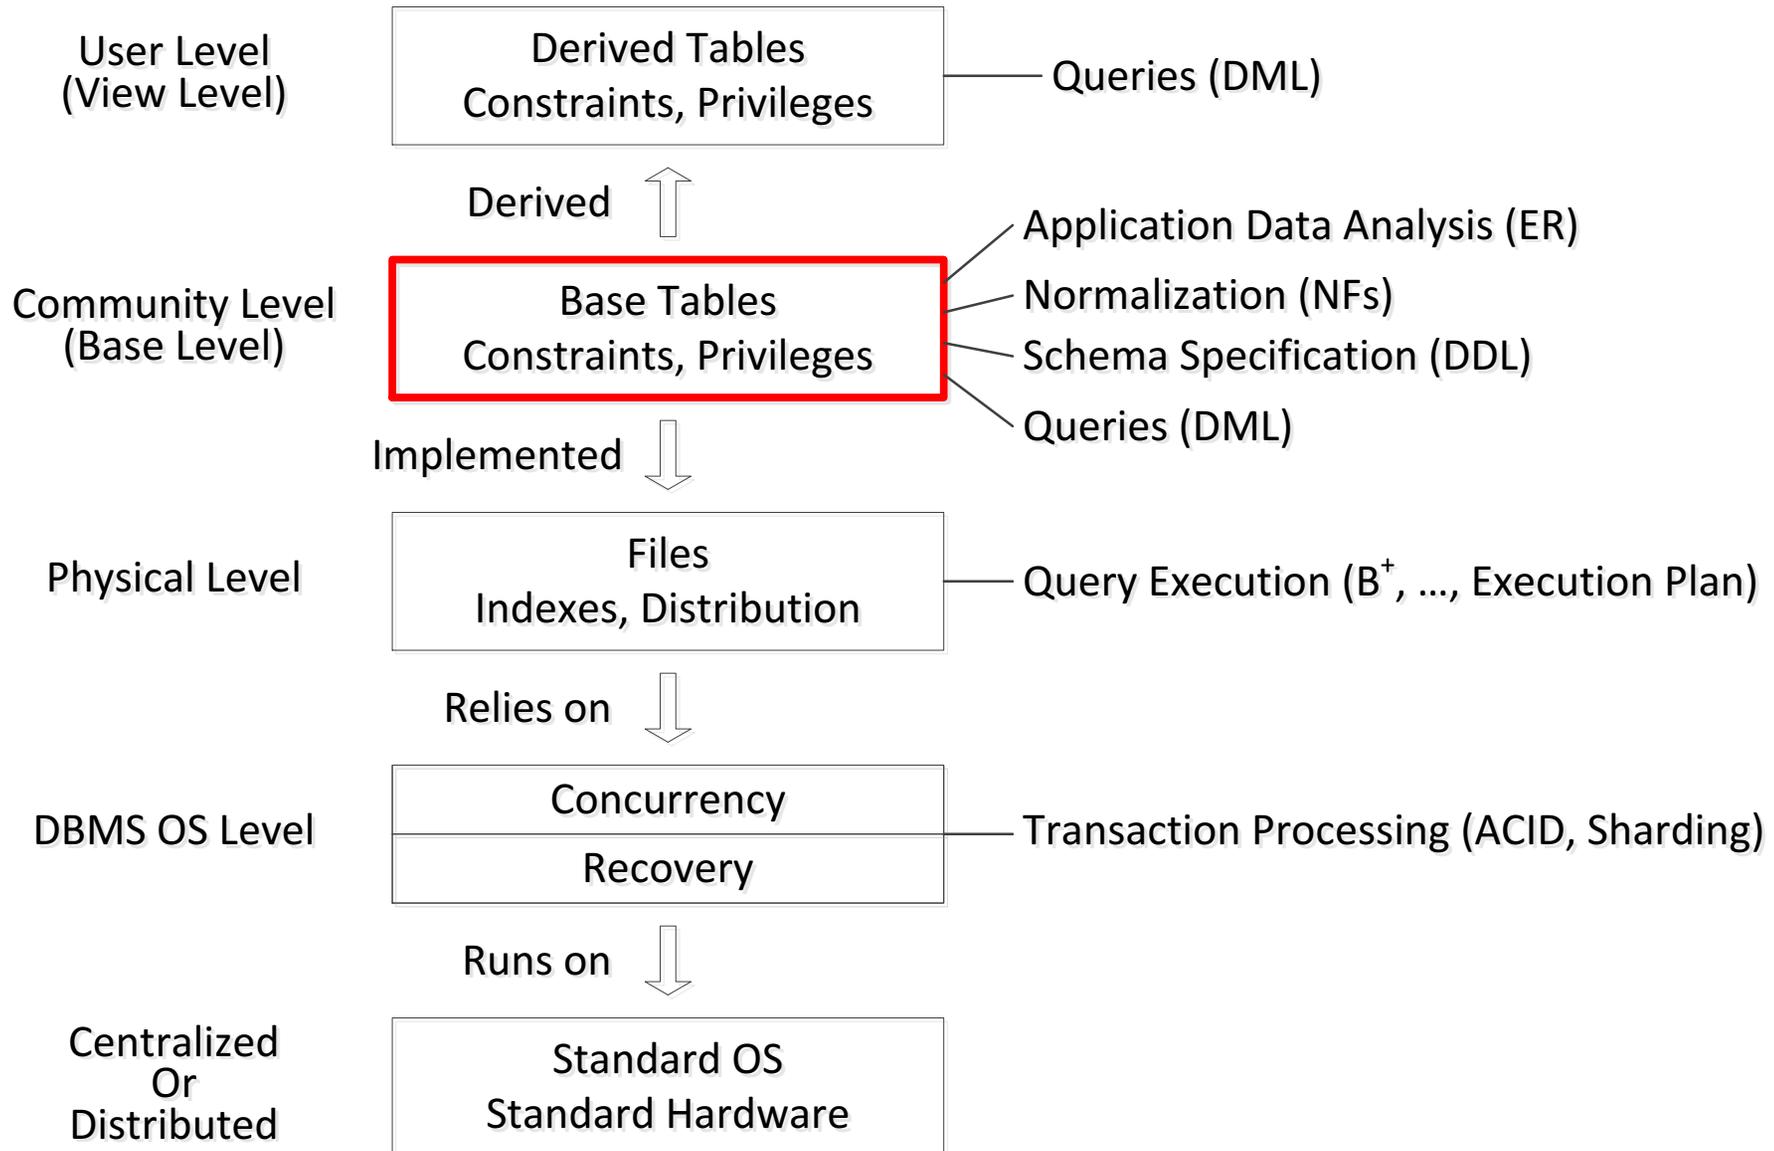


Unit 2
***Modeling the Information of an Enterprise Using
Chen's Entity/Relationship Model and Diagrams***

ER Diagrams in Context



Purpose Of ER Model And Basic Concepts

- ◆ ***Entity/relationship (ER) model*** provides a common, informal, and convenient method for communication between application end users (customers) and the database designers to model the information's structure
- ◆ This is a preliminary stage towards defining the database using a formal model, such as the relational model, to be described later
- ◆ The ER model, frequently employs ***ER diagrams***, which are pictorial descriptions to visualize information's structure
- ◆ ER models are both ***simple*** and ***powerful***

Purpose Of ER Model And Basic Concepts

- ◆ There are three basic concepts appearing in the original ER model, which has since been extended
 - We will present the model from more simple to more complex concepts, with examples on the way
- ◆ We will go beyond the original ER model, and cover most of the ***Enhanced ER*** model
- ◆ While the ER model's ***concepts are standard***, there are several ***varieties of pictorial representations*** of ER diagrams
 - We will focus on one of them: ***Chen's*** notation
 - We will also cover ***Crow's foot*** notation in the context of the Visio tool
 - Others are simple variations, so if we understand the above, we can easily understand all of them
- ◆ You can look at some examples at:
http://en.wikipedia.org/wiki/Entity-relationship_model

Basic Concepts

- ◆ The three basic concepts are (elaborated on very soon):
- ◆ **Entity**. This is an “object.” Cannot be defined even close to a formal way. Examples:
 - Bob
 - Boston
 - The country whose capital is Paris
There is only one such country so it is completely specified
- ◆ **Relationship**. Entities participate in relationships with each other. Examples:
 - Alice and Boston are in relationship Likes (Alice likes Boston)
 - Bob and Atlanta are not in this relationship
- ◆ **Attribute** (property). Examples:
 - Age is an attribute of persons
 - Size is an attribute of cities

Entity And Entity Set

- ◆ ***Entity*** is a “thing” that is distinguished from others in our application
 - Example: Alice
- ◆ All entities of the same “type” form an ***entity set***; we use the term “type” informally
 - Example: Person (actually a set of persons). Alice is an entity in this entity set
- ◆ What type is a little tricky sometimes
- ◆ Example. Do we partition people by sex or not?
 - Sometimes makes sense (gave birth)

This allows better enforcement of constraints. You could “automatically” make sure that only entities in the set of women, but not in the set of men can give birth
 - Sometimes not (employment)

Entity And Entity Set

- ◆ Example. When we say “the set of all Boeing airplanes,” is this
 - The set of all models appearing in Boeing’s catalog (abstract objects), or
 - The set of airplanes that Boeing manufactured (concrete objects)
- ◆ We may be interested in both and have two entity sets that are somehow related
- ◆ We will frequently use the term “entity” while actually referring to entity sets, unless this causes confusion

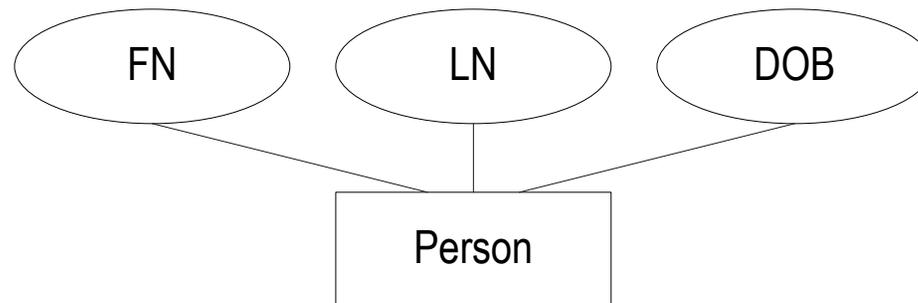
Entity And Entity Set

- ◆ Pictorially, an entity set is denoted by a **rectangle** with its type written inside
- ◆ By convention, singular noun, though we may not adhere to this convention if not adhering to it makes things clearer
- ◆ By convention, capitalized, or all capitals, if acronym



Attribute

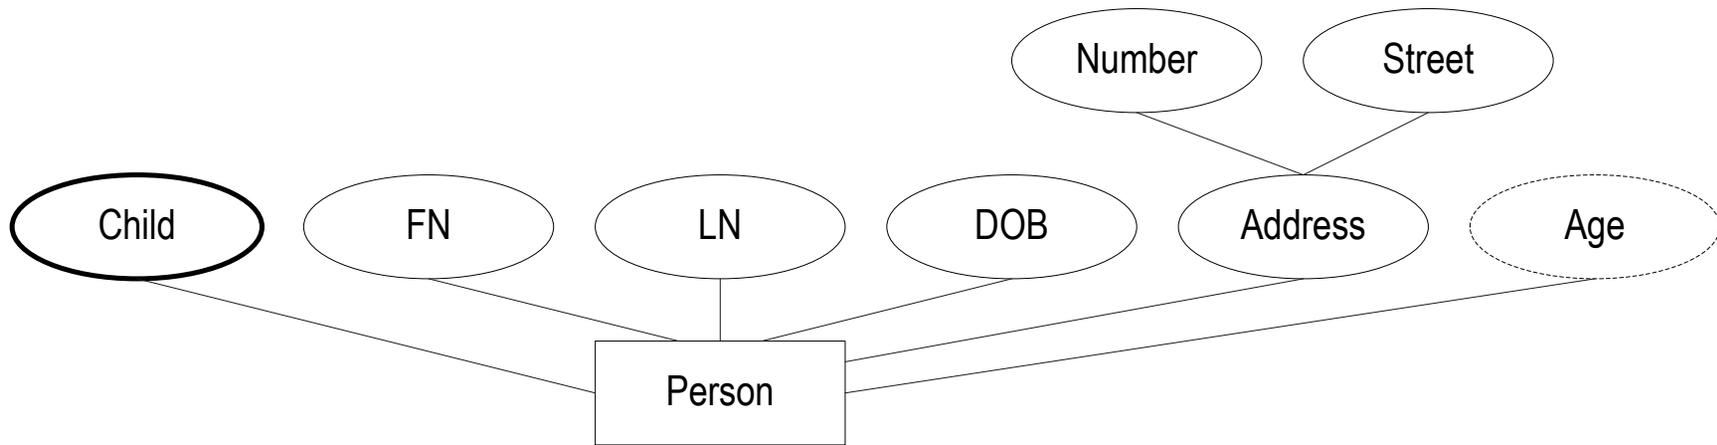
- ◆ An entity may have (and in general has) a set of zero or more **attributes**, which are some properties
- ◆ Each attribute is drawn from some domain (such as integers) possibly augmented by **NULL** (more about NULLs later)
- ◆ All entities in an entity set have the same set of properties, though not generally with the same values
- ◆ Attributes of an entity are written in ellipses (for now solid lines) connected to the entity
 - Example: FN: “First Name.” LN: “Last Name.” DOB: “Date of Birth.”



Attribute

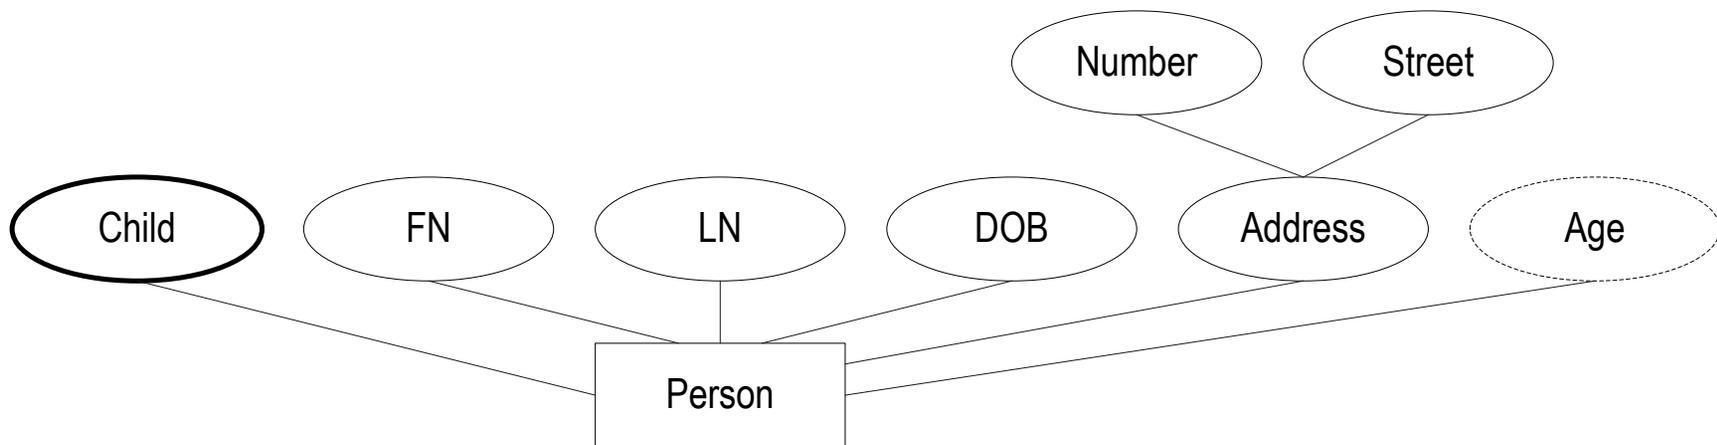
◆ Attributes can be

- **Base** (such as DOB); or **derived** denoted by dashed ellipses (such as Age, derived from DOB and the current date)
- **Simple** (such as DOB); or **composite** having their component attributes attached to them (such as Address, when we think of it explicitly as consisting of street and number and restricting ourselves to one city only)
- **Singlevalued** (such as DOB); or **multivalued** with unspecified in advance number of values denoted by thick-lined ellipses (such as Child; a person may have any number of children; we do not consider children as persons in this example, this means that they are not elements of the entity set Person, just attributes of elements of this set)



Attribute

- ◆ To have a simple example of a person with attributes
 - Child: Bob
 - Child: Carol
 - FN: Alice
 - LN: Xie
 - DOB: 1980-01-01
 - Address.Number: 100
 - Address.Street: Mercer
 - Age: Current Date minus DOB specified in years (rounded down)

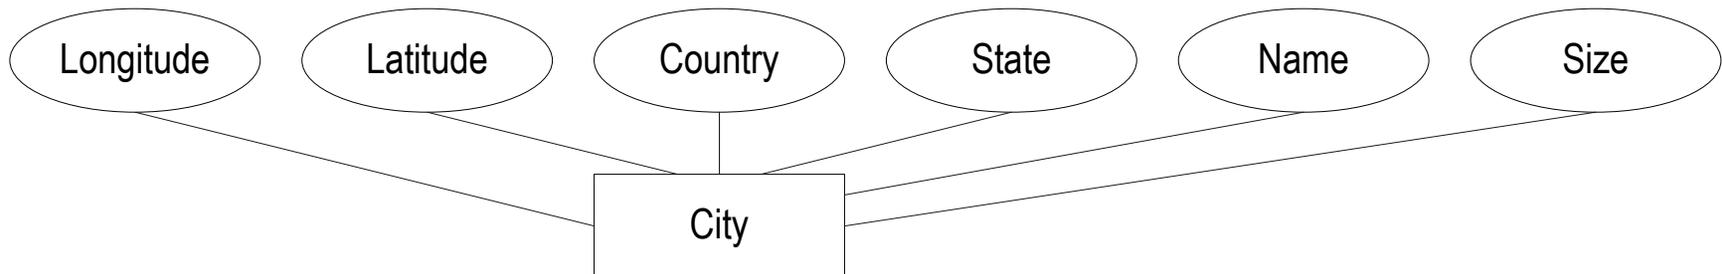


Sets, Subsets, and Supersets

- ◆ Relations subset and superset are defined among sets
- ◆ It is analogous to
- ◆ Let us review by an example of three sets
 - $A = \{2,5,6\}$
 - $B = \{1,2,5,6,8\}$
 - $C = \{2,5,6\}$
- ◆ Then we have
 - $A \subset B$ and A is a subset of B
and A is a proper subset, actually is not all of B ; $A \subsetneq B$
 - $A \subset C$ and A is a subset of C
and A is not a proper subset, actually is equal to C ; $A = C$
- ◆ Caution: sometimes \subseteq is used to denote what we denote by \subset

Keys

- ◆ Most of the times, some subset (proper or not) of the attributes of an entity has the property that **two different entities in an entity set must differ on the values of these attributes**
- ◆ This must hold for all conceivable entities in our database
- ◆ Such a set of attributes is called a **superkey** (“weak” superset of a key: either proper superset or equal)
- ◆ A minimal superkey is called a **key** (sometimes called a **candidate key**).
 - This means that no proper subset of it is itself a superkey

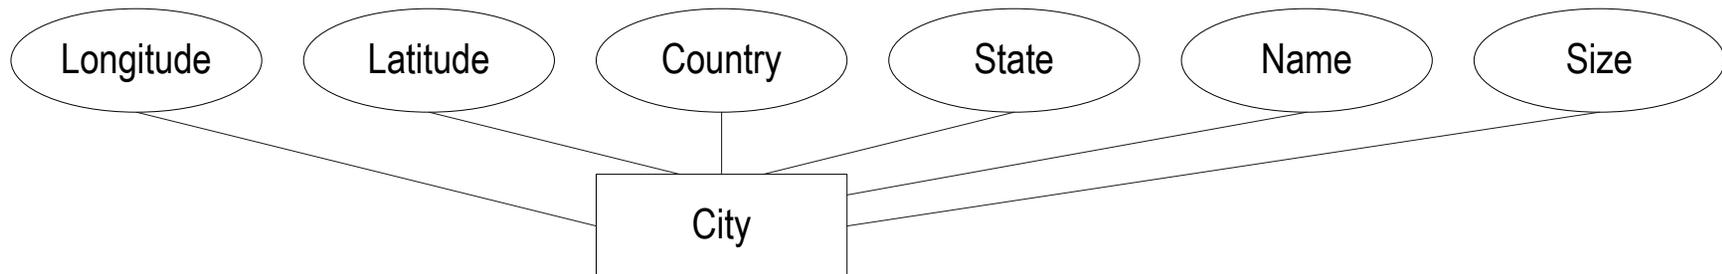


Keys

- ◆ Informally: superkey values can identify an individual entity but there may be unnecessary attributes
- ◆ Informally: key value can identify an individual entity but there are no unnecessary attributes
- ◆ Example: Social Security Number + Last Name form a superkey, which is not a key as Social Security Number is enough to identify a person

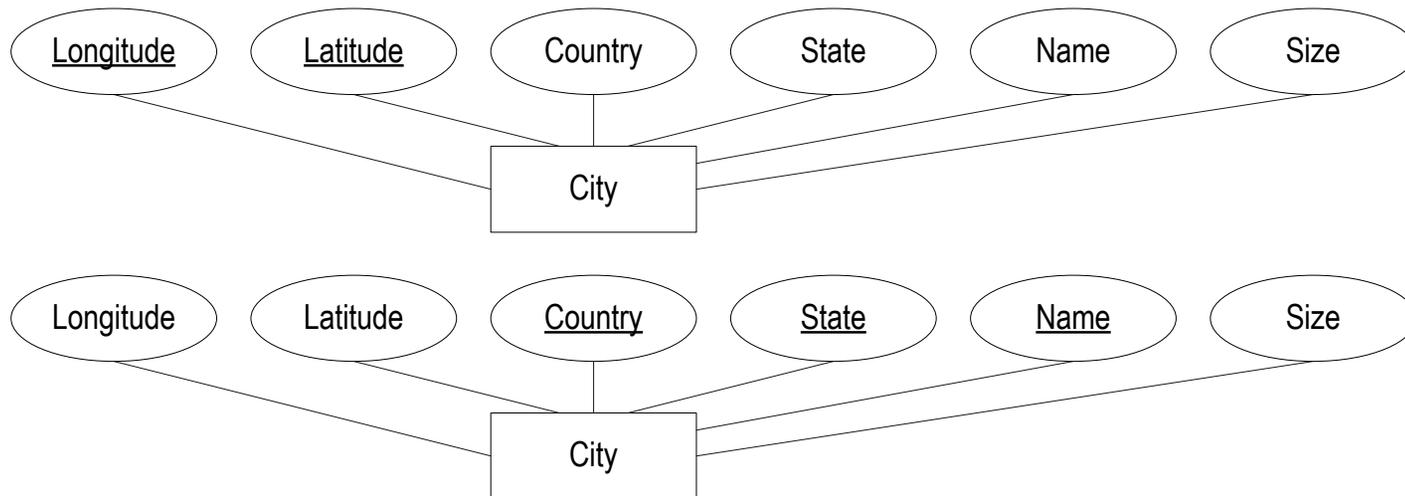
Keys

- ◆ In our example:
 - Longitude and Latitude (their values) identify (at most) one City, but only Longitude or only Latitude do not
 - (Longitude, Latitude) form a superkey, which is also a key
 - (Longitude, Latitude, Size, Name) form a superkey, which is not a key, because Size and Name are superfluous
 - (Country, State, Name) form another key (and also a superkey, as every key is a superkey)
- ◆ For simplicity, we assume that every country is divided into states and within a state the city name is unique



Primary Keys

- ◆ If an entity set has one or more keys, one of them (no formal rule which one) is chosen as the **primary key**
- ◆ In SQL the other keys, loosely speaking, are referred to using the keyword **UNIQUE**
- ◆ In the ER diagram, the attributes of the primary key are underlined
- ◆ So in our example, **one** of the two below:



Relationship

- ◆ Several entity sets (one or more) can participate in a *relationship*
- ◆ Relationships are denoted by diamonds, to which the participating entities are “attached”
- ◆ A relationship could be binary, ternary,
- ◆ By convention, a capitalized verb in third person singular (e.g., Likes), though we may not adhere to this convention if not adhering to it makes things clearer

Relationship

- ◆ We will have some examples of relationships
- ◆ We will use three entity sets, with entities (and their attributes) in those entity sets listed below

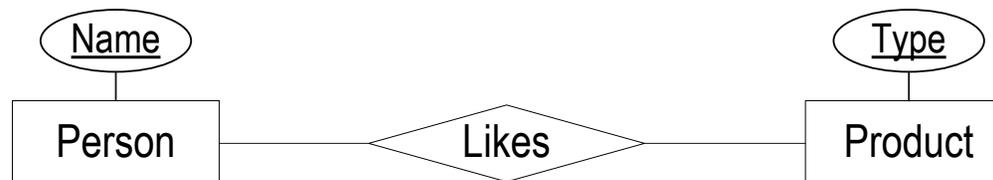
Person	Name
	Chee
	Lakshmi
	Marsha
	Michael
	Jinyang

Vendor	Company
	IBM
	Apple
	Dell
	HP

Product	Type
	computer
	monitor
	printer

Binary Relationship

- ◆ Let's look at Likes, listing all pairs of (x,y) where person x Likes product y , and the associated ER diagram
- ◆ First listing the relationship informally (we omit article “a”):
 - Chee likes computer
 - Chee likes monitor
 - Lakshmi likes computer
 - Marsha likes computer
- ◆ Note
 - Not every person has to Like a product
 - Not every product has to have a person who Likes it (informally, be Liked)
 - A person can Like many products
 - A product can have many person each of whom Likes it



Relationships

- ◆ Formally we say that R is a **relationship** among (not necessarily distinct) entity sets E_1, E_2, \dots, E_n if and only if R is a subset of $E_1 \times E_2 \times \dots \times E_n$ (Cartesian product)
- ◆ In our example above:
 - $n = 2$
 - $E_1 = \{\text{Chee, Lakshmi, Marsha, Michael, Jinyang}\}$
 - $E_2 = \{\text{computer, monitor, printer}\}$
 - $E_1 \times E_2 = \{ (\text{Chee,computer}), (\text{Chee,monitor}), (\text{Chee,printer}), (\text{Lakshmi,computer}), (\text{Lakshmi,monitor}), (\text{Lakshmi,printer}), (\text{Marsha,computer}), (\text{Marsha,monitor}), (\text{Marsha,printer}), (\text{Michael,computer}), (\text{Michael,monitor}), (\text{Michael,printer}), (\text{Jinyang,computer}), (\text{Jinyang,monitor}), (\text{Jinyang,printer}) \}$
 - $R = \{ (\text{Chee,computer}), (\text{Chee,monitor}), (\text{Lakshmi,computer}), (\text{Marsha,monitor}) \}$
- ◆ R is a set (unordered, as every set) of ordered tuples, or sequences (here of length two, that is pairs)

Relationships

- ◆ Let us elaborate
- ◆ $E_1 \times E_2$ was the “universe”
 - It listed all possible pairs of a person liking a product
- ◆ At every instance of time, in general only some of this pairs corresponded to the “actual state of the universe”; R was the set of such pairs

Important Digression

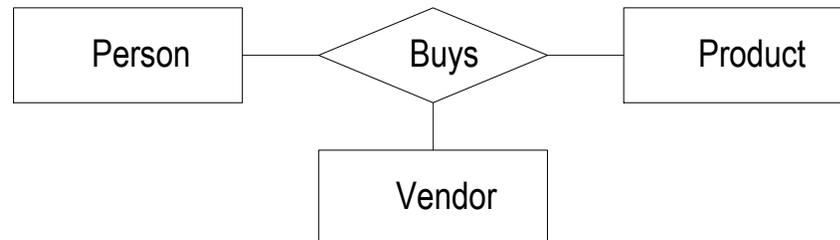
- ◆ Ultimately, we will store (most) relationships as tables
- ◆ So, our example for Likes could be

Likes	Name	Type
	Chee	Computer
	Chee	Monitor
	Lakshmi	Computer
	Marsha	Monitor

- ◆ Where we identify the “participating” entities using their primary keys

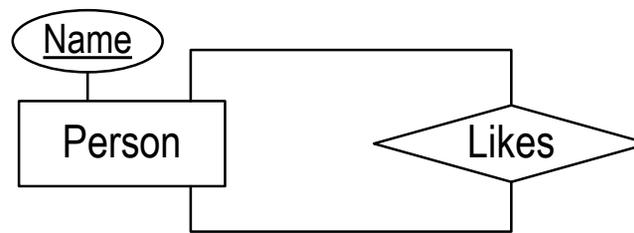
Ternary Relationship

- ◆ Let's look at Buys listing all tuples of (x,y,z) where person x Buys product y from vendor z
- ◆ Let us just state it informally:
 - Chee buys computer from IBM
 - Chee buys computer from Dell
 - Lakshmi buys computer from Dell
 - Lakshmi buys monitor from Apple
 - Chee buys monitor from IBM
 - Marsha buys computer from IBM
 - Marsha buys monitor from Dell



Relationship With Nondistinct Entity Sets

- ◆ Let's look at Likes, listing all pairs of (x,y) where person x Likes person y
- ◆ Let us just state it informally
 - Chee likes Lakshmi
 - Chee likes Marsha
 - Lakshmi likes Marsha
 - Lakshmi likes Michael
 - Lakshmi likes Lakshmi
 - Marsha likes Lakshmi
- ◆ Note that pairs must be ordered to properly specify the relationship, Chee likes Lakshmi, but Lakshmi does not like Chee



Relationship With Nondistinct Entity Sets

- ◆ Again:
 - Chee likes Lakshmi
 - Chee likes Marsha
 - Lakshmi likes Marsha
 - Lakshmi likes Michael
 - Lakshmi likes Lakshmi
 - Marsha likes Lakshmi

- ◆ Formally Likes is a subset of the Cartesian product $\text{Person} \times \text{Person}$, which is the set of all ordered pairs of the form (person, person)

- ◆ Likes is the set $\{ (\text{Chee}, \text{Lakshmi}), (\text{Chee}, \text{Marsha}), (\text{Lakshmi}, \text{Marsha}), (\text{Lakshmi}, \text{Michael}), (\text{Lakshmi}, \text{Lakshmi}), (\text{Marsha}, \text{Lakshmi}) \}$

- ◆ Likes is an arbitrary directed graph in which persons serve as vertices and arcs specify who likes whom

Important Digression

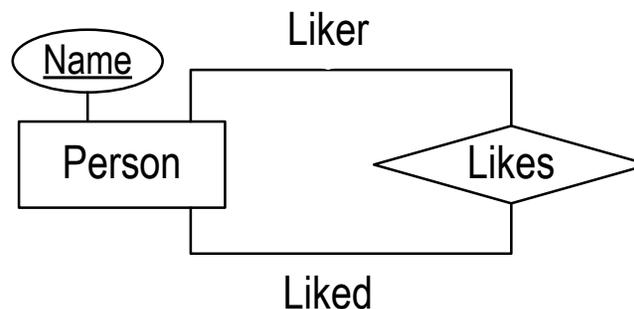
- ◆ Ultimately, we will store (most) relationships as tables
- ◆ So, our example for Likes could be

Likes	Name	Name
	Chee	Lakshmi
	Chee	Marsha
	Lakshmi	Marsha
	Lakshmi	Michael
	Lakshmi	Lakshmi
	Marsha	Lakshmi

- ◆ Where we identify the “participating” entities using their primary keys
- ◆ But it is difficult to see (unless we keep track of columns order) whether Lakshmi Likes Michael or Michael Likes Lakshmi

Relationship With Nondistinct Entity Sets

- ◆ Frequently it is useful to give **roles** to the participating entities, when, as here, they are drawn from the same entity set.
- ◆ So, we may say that if Chee likes Lakshmi, then Chee is the “Liker” and Lakshmi is the “Liked”
- ◆ Roles are explicitly listed in the diagram, but the semantics of they mean cannot be deduced from looking at the diagram only



Important Digression

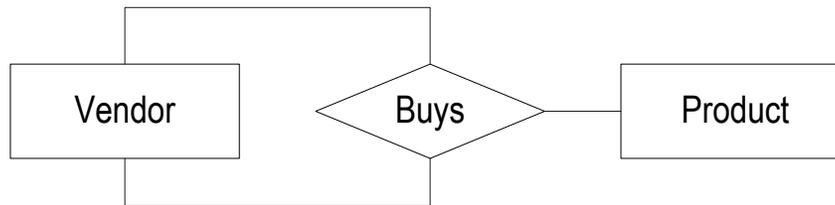
- ◆ Ultimately, we will store (most) relationships as tables
- ◆ So, our example for Likes could be

Likes	Liker	Liked
	Chee	Lakshmi
	Chee	Marsha
	Lakshmi	Marsha
	Lakshmi	Michael
	Lakshmi	Lakshmi
	Marsha	Lakshmi

- ◆ Where we identify the “participating” entities using their primary keys ***but we rename them using roles***
- ◆ So we do not need to keep track of columns order and we know that Lakshmi Likes Michael and Michael does not Like(s) Lakshmi, though we still do not know what “Likes” really means

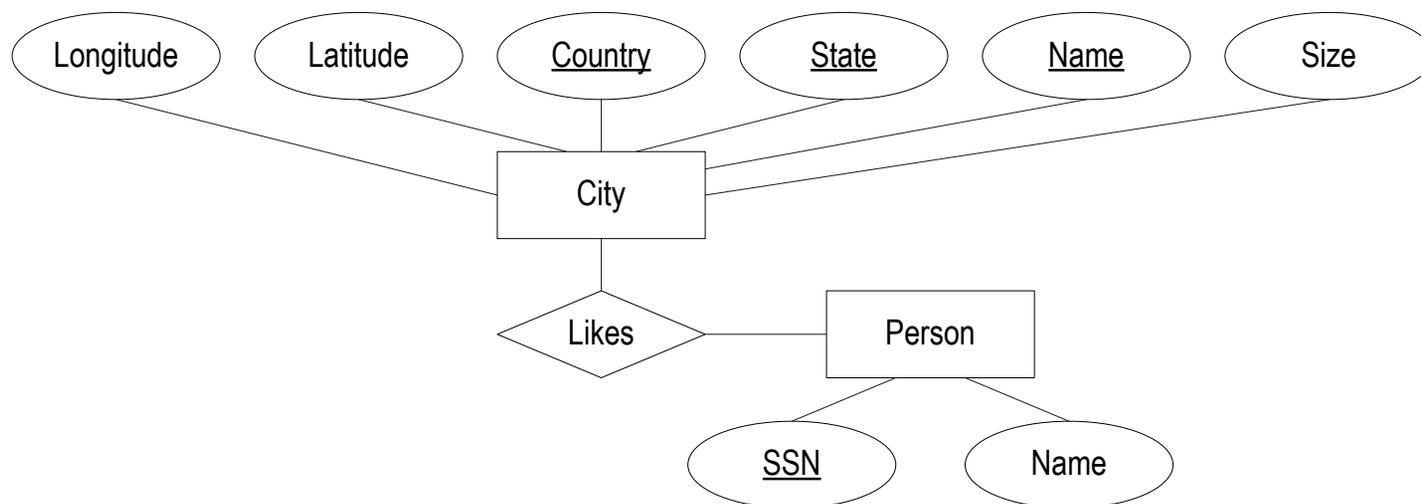
Relationship With Nondistinct Entity Sets

- ◆ Consider Buys, listing all triples of the form (x,y,z) where vendor x Buys product y from vendor z
- ◆ A typical tuple might be $(Dell,printer,HP)$, meaning that Dell buys a printer from HP



ER Diagrams

- ◆ To show which entities participate in which relationships, and which attributes participate in which entities, we draw line segments between:
 - Entities and relationships they participate in
 - Attributes and entities they belong to
- ◆ We also underline the attributes of the primary key for each entity that has a primary key
- ◆ Below is a simple ER diagram (with a simpler Person than we had before):



Further Refinements To The ER Model

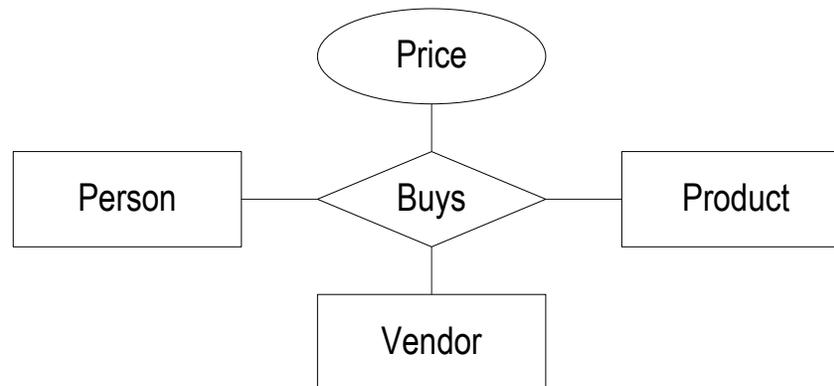
- ◆ We will present, in steps, further refinements to the model and associated diagrams
- ◆ The previous modeling concepts and the ones that follow are needed for producing a data base design that models a given application well
- ◆ We will then put it together in a larger comprehensive example

Relationship With Attributes

- ◆ Consider relationship Buys among Person, Vendor, and Product
- ◆ We want to specify that a person Buys a product from a vendor at a specific price
- ◆ Price is not
 - A property of a vendor, because different products may be sold by the same vendor at different prices
 - A property of a product, because different vendors may sell the same product at different prices
 - A property of a person, because different products may be bought by the same person at different prices

Relationship With Attributes

- ◆ So Price is really an attribute of the relationship Buys
- ◆ For each tuple (person, product, vendor) there is a value of price



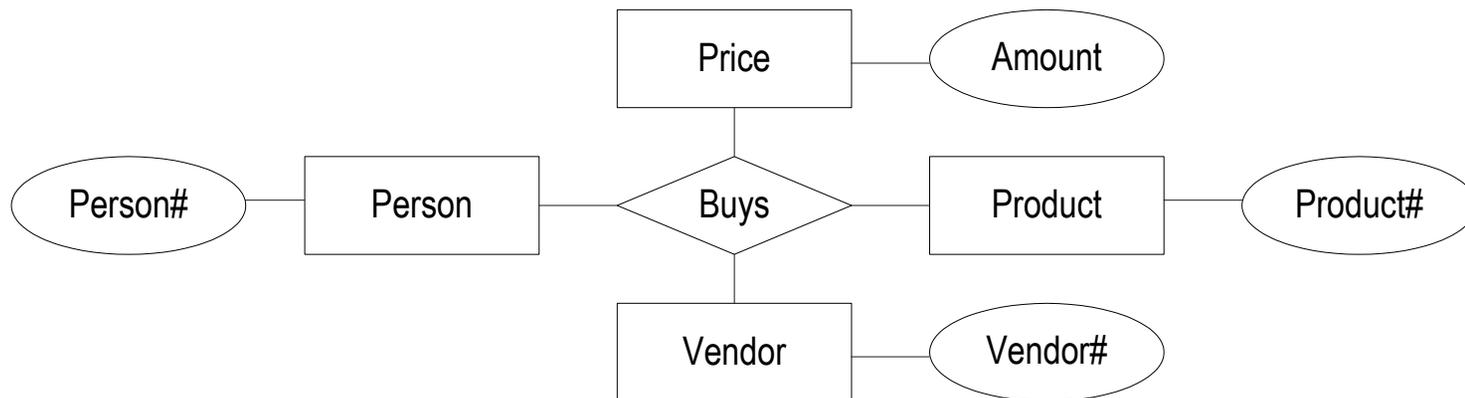
Entity Versus Attribute

- ◆ Entities can model situations that attributes cannot model naturally
- ◆ Entities can
 - Participate in relationships
 - Have attributes
- ◆ Attributes cannot do any of these

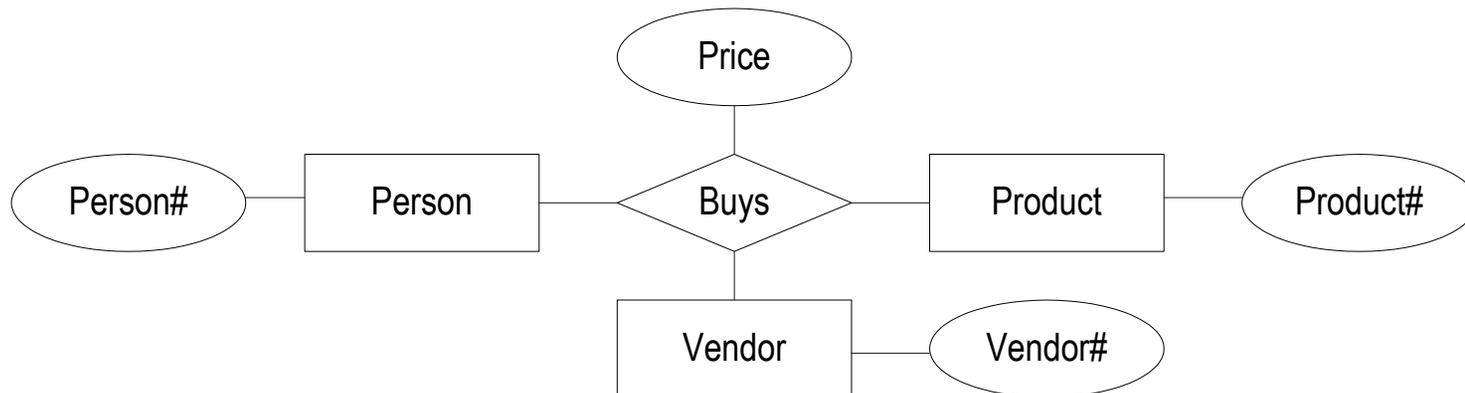
- ◆ Let us look at a “fleshed out example” for possible alternative modeling of Buys

Other Choices For Modeling Buys

- ◆ Price is just the actual amount, the number in \$'s
- ◆ So there likely is no reason to make it an entity as we have below

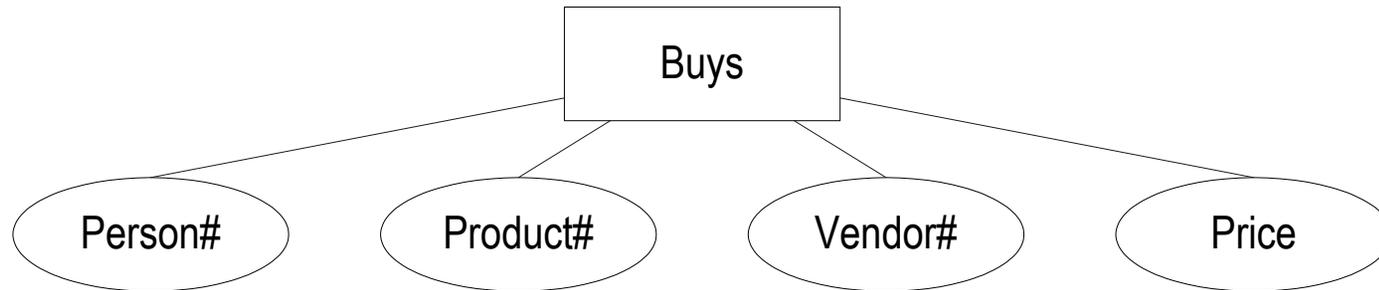


- ◆ We should probably have (as we had earlier less fleshed out)



Other Choices For Modeling Buys

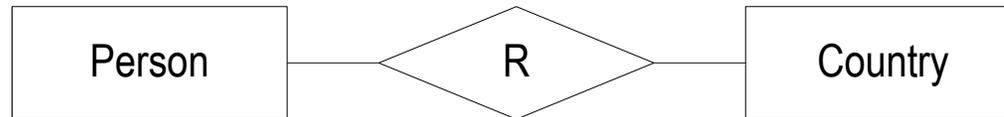
- ◆ Or should we just have this?



- ◆ Not if we want to model something about a person, such as the date of birth of a person or whom a person likes
- ◆ These require a person to have an attribute (date of birth) and enter into a relationship (with other persons)
- ◆ And we cannot model this situation if person is an attribute of Buy
- ◆ Similarly, for product and vendor

Binary Relationships And Their Functionality

- ◆ Consider a relationship R between two entity sets A, B.
- ◆ We will look at examples where A is the set of persons and B is the set of all countries



- ◆ We will be making some simple assumptions about persons and countries, which we list when relevant

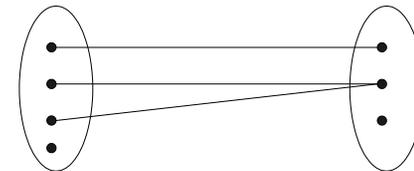
Binary Relationships And Their Functionality

- ◆ Relationship R is called **many to one** from A to B if and only if for each element of A there exists at most one element of B related to it

- Example: R is Born (in)

Each person was born in at most one country (maybe not in a country but on a ship in the middle of an ocean)

Maybe nobody was born in some country as it has just been established

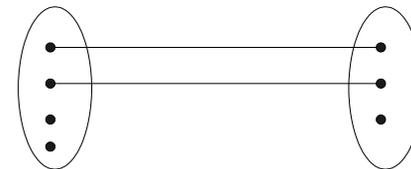


- ◆ The picture on the right describes the universe of four persons and three countries, with lines indicating which person was born in which country

- We will have similar diagrams for other examples

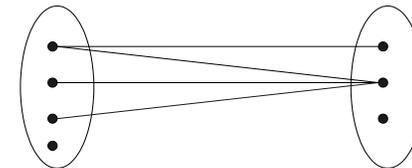
Binary Relationships And Their Functionality

- ◆ The relationship R is called **one to one** between A and B if and only if for each element of A there exists at most one element of B related to it and for each element of B there exists at most one element of A related to it
 - Example: R is Heads
Each Person is a Head (President, Queen, etc.) of at most one country
Each country has at most one head (maybe the queen died and it is not clear who will be the monarch next)
- ◆ In other words, R is one to one, if and only if
 - R is many to one from A to B, and
 - R is many to one from B to A



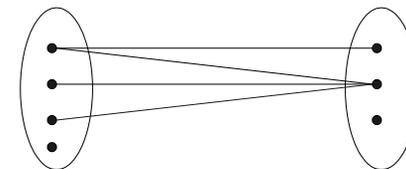
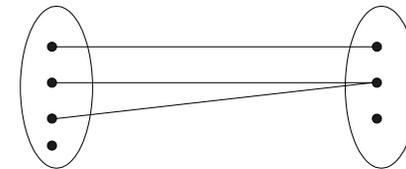
Binary Relationships And Their Functionality

- ◆ The relationship is called **many to many** between A and B, if it is not many to one from A to B and it is not many to one from B to A
 - Example: R is “likes”



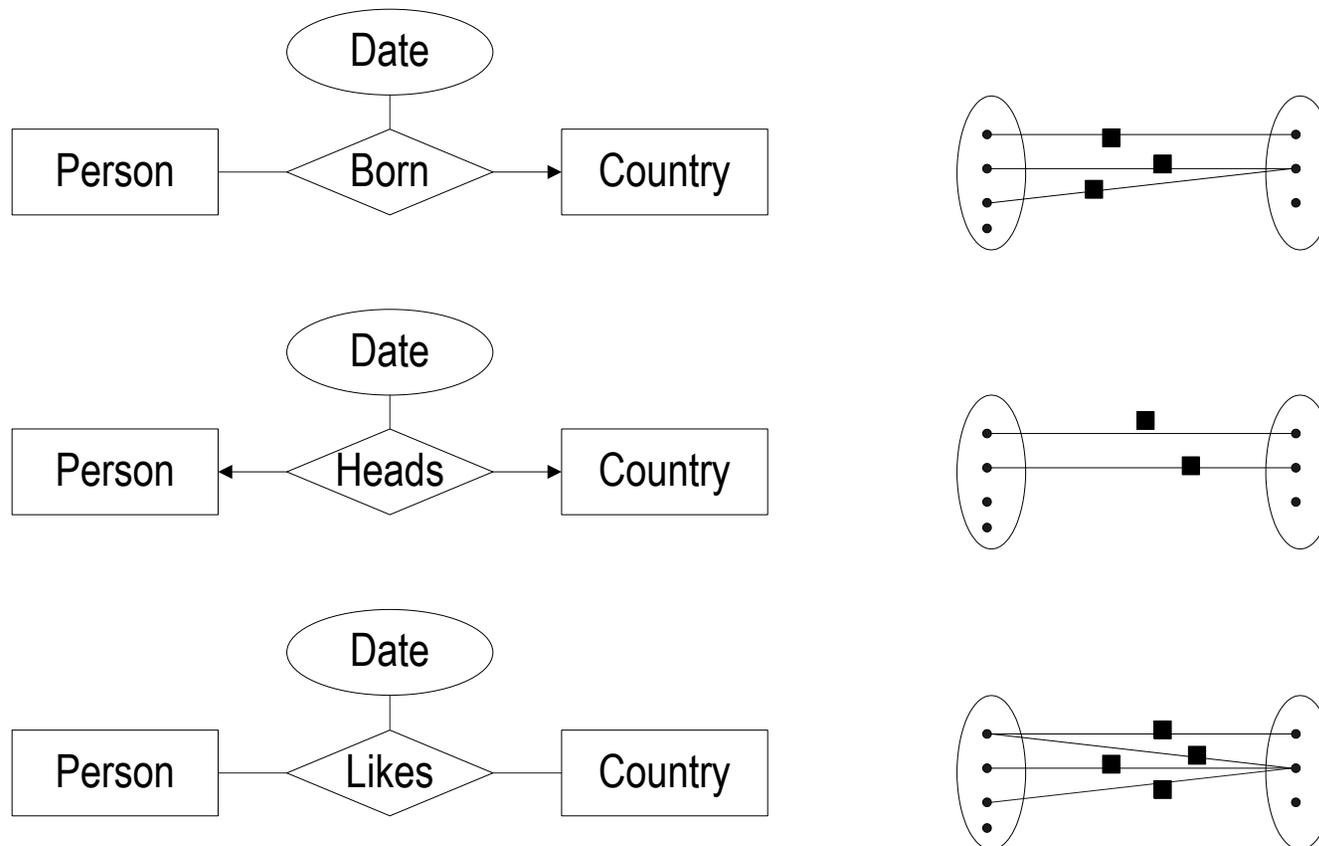
Binary Relationships And Their Functionality

- ◆ We have in effect considered the concepts of partial functions of one variable.
 - The first two examples were *partial functions*
 - The last example was not a function
- ◆ Pictorially, functionality for binary relationships can be shown by drawing an arc head in the direction to the “one”



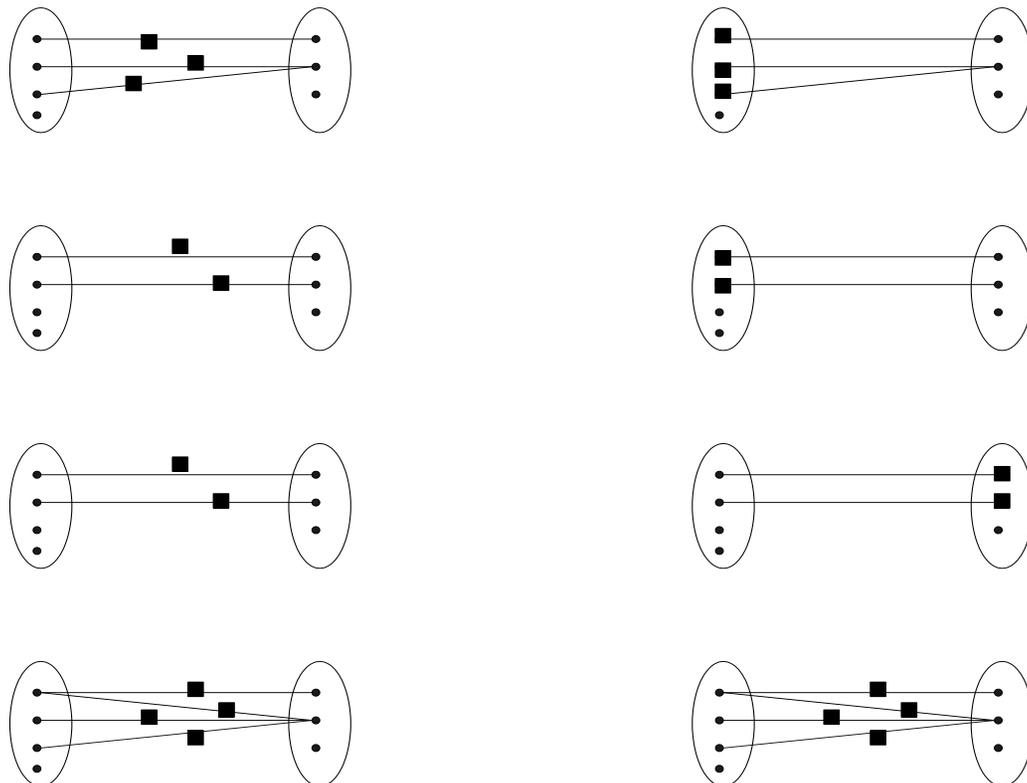
Binary Relationships And Their Functionality

- ◆ How about properties of the relationship?
- ◆ Date: when a person and a country in a relationship first entered into the relationship (marked also with black square)



Binary Relationships And Their Functionality

- ◆ Can make Date in some cases the property of an entity
 - “Slide” the Date to the Person, but not the Country
 - “Slide” the Date to either the Person or the Country (but not for both, as this would be redundant)
- ◆ Cannot “slide” the Date to either “Liker” or “Liked”
- ◆ Can “slide” if no two squares end up in the same entity

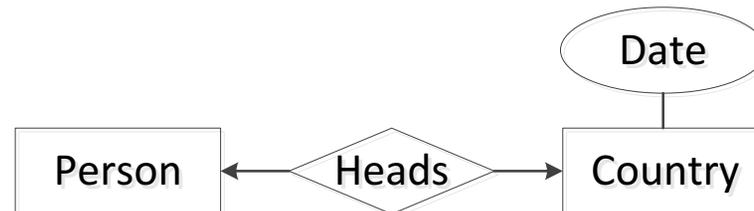
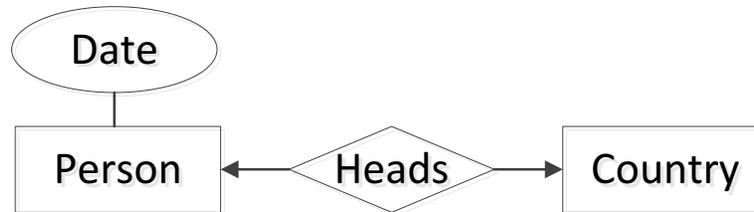
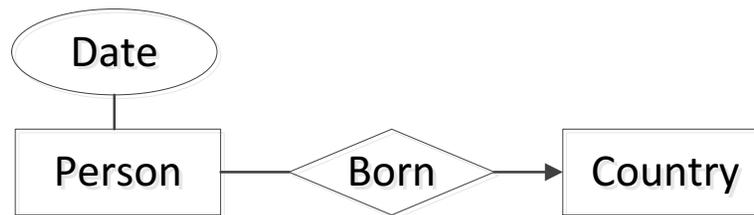


Binary Relationships And Their Functionality

- ◆ This can be done if the relationship is many-to-one
 - ◆ Then, the property of the relationship can be attributed to the “many” side
-
- ◆ This can be done if the relationship is one-to-one
 - ◆ Then a property of the relationship can be “attributed” to any of the two sides

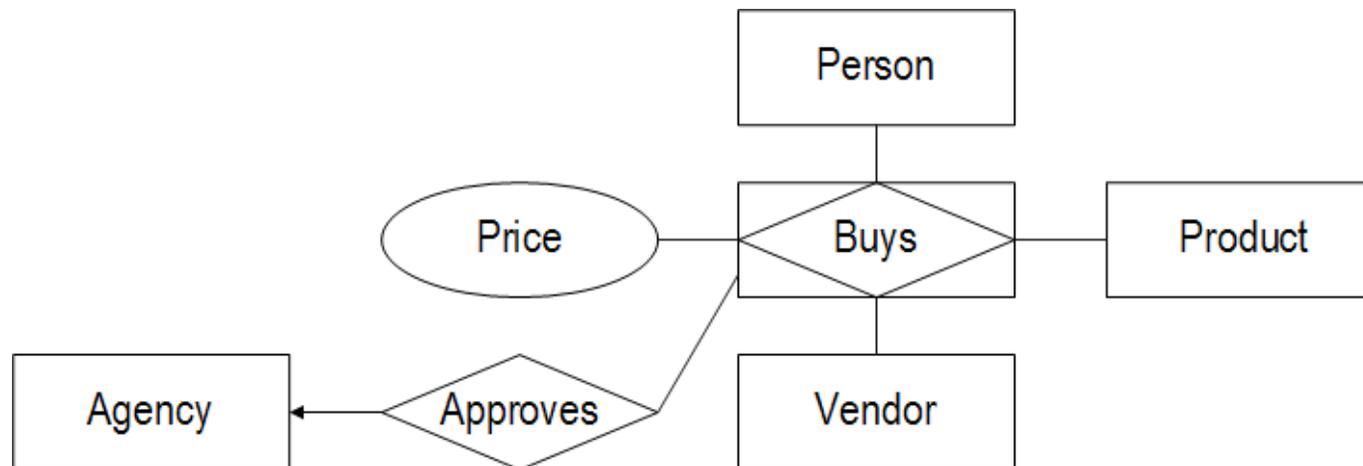
Alternate Designs

- ◆ Entities “inheriting” attributes of relationships when the relationships are not many to many



Aggregation: Relationships As Entities

- ◆ It is sometimes natural to consider relationships as if they were entities.
- ◆ This will allow us to let relationships participate in other “higher order” relationships
- ◆ Here each “contract” needs to be approved by (at most) one agency
- ◆ Relationship is “made into” an entity by putting it into a **rectangle**; note that the edge between Buys and Approves touches the Buys rectangle but not the Buys diamond, to make sure we are not confused



Strong And Weak Entities

- ◆ We have two entity sets
 - Man
 - Woman
- ◆ Woman has a single attribute, SSN
- ◆ Let us defer discussion of attributes of Man
- ◆ A woman has 5 sons, the among them John and Richard, neither of the two is her eldest son and she writes the following in her will:

My SSN is 123-45-6789 and I leave \$100 to my eldest son and \$200 to my son John and \$300 to my son Richard ...

- ◆ How do we identify these 3 men?

Strong And Weak Entities

- ◆ A ***strong entity*** (set): Its elements can be identified by the values of their attributes, that is, it has a (primary) key made of its attributes

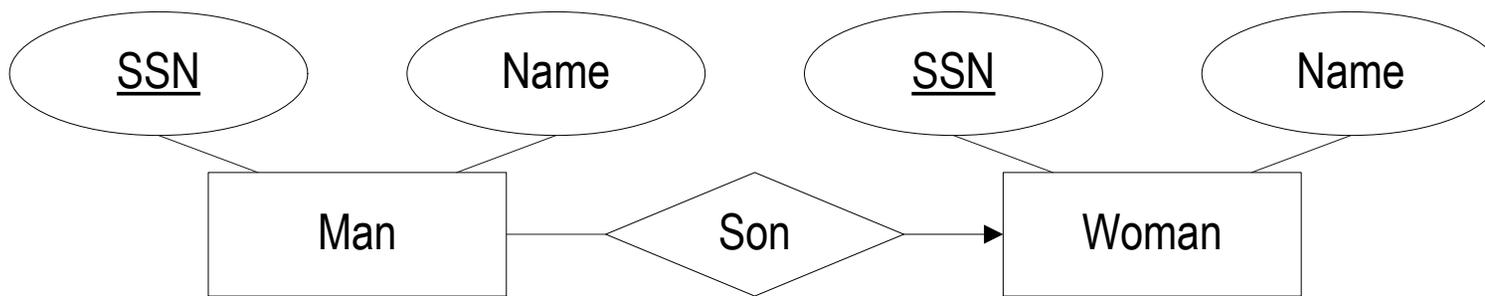
Tacitly, we assumed only such entities so far

- ◆ A ***weak entity*** (set): Its elements cannot be identified by the values of their attributes: there is no primary key made from its own attributes

Such entities can be identified by a combination of their attributes and the relationship they have with another entity set

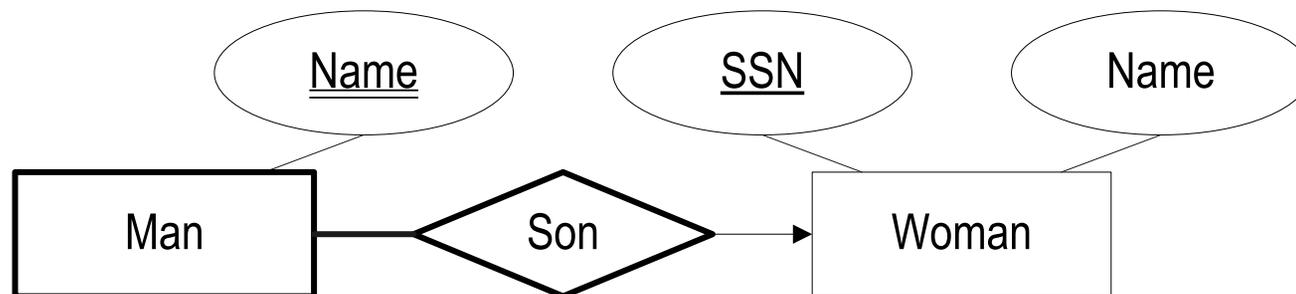
Man As A Strong Entity

- ◆ Most entities are **strong**: a specific entity can be distinguished from other entities based on the values of its attributes
- ◆ We assume that every person has his/her own SSN
- ◆ Woman is a strong entity as we can identify a specific woman based on her attributes. She has a primary key: her own SSN
- ◆ Man is a strong entity as we can identify a specific man based on his attributes. He has a primary key: his own SSN



Man As A Weak Entity

- ◆ We assume that women are given SSNs
- ◆ Men are not given SSNs; they have first names only, but for each we know who the mother is (that is, we know the SSN of the man's mother)
- ◆ Man is a **weak** entity as we cannot identify a specific man based on his own attributes and **this is indicated by thick lines around it**
- ◆ Many women could have a son named Bob, so there are many men named Bob
- ◆ However, if a woman never gives a specific name to more than one of her sons, a man can be identified by his name and by his mother's SSN



Man As A Weak Entity

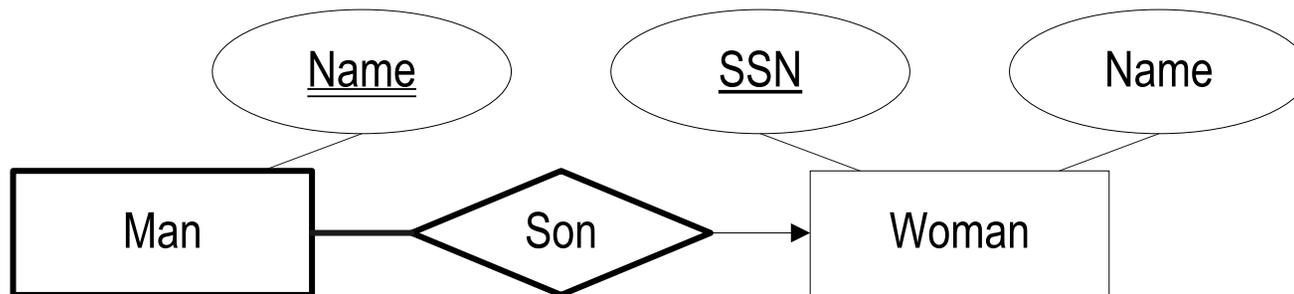
- ◆ We could have the following situation of two mothers: one with two sons, and one with three sons, when we gave people also heights in inches (just to have additional attributes that are not necessary for identification)

- ◆ SSN: 070-43-1234, height: 65
 - Name: Bob, height 35
 - Name: Michael, height 35

- ◆ SSN: 056-35-4321, height 68
 - Name: Bob, height 35
 - Name: Davi, height 45
 - Name: Vijay, height 74

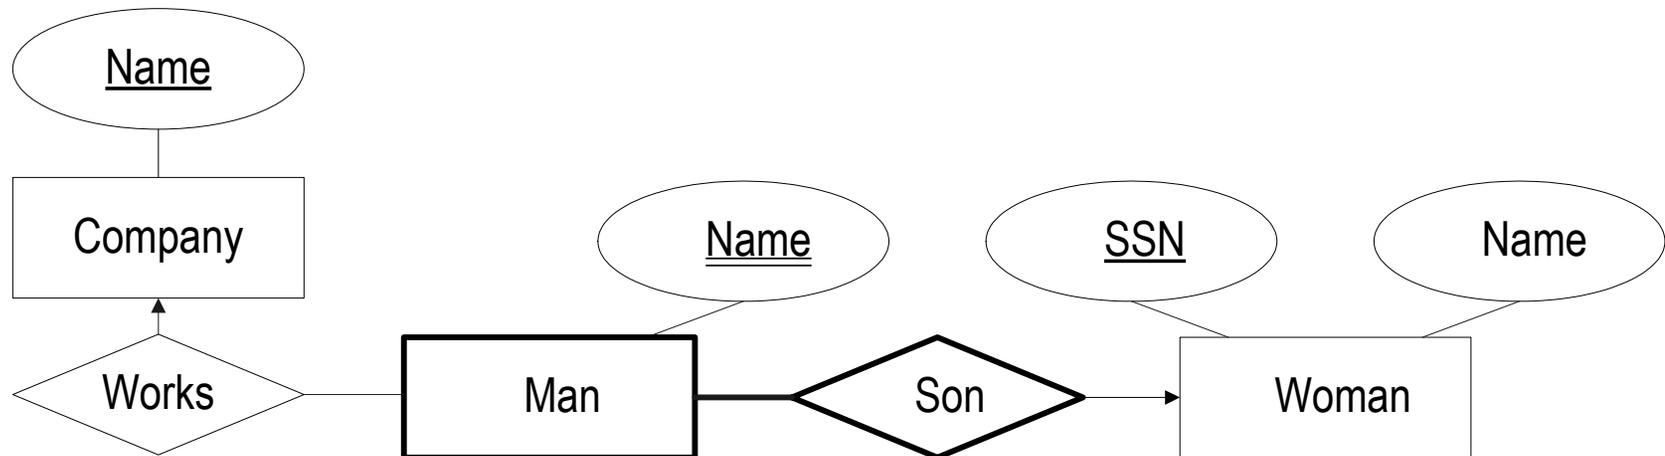
Man As A Weak Entity

- ◆ Assuming that a woman does not have more than one son with a specific name
- ◆ Name becomes a **discriminant**
- ◆ Man can be identified by the combination of:
 - The Woman to whom he is related under the Son relation. This is indicated by thick lines around Son (it is weak). Thick line connecting Man to Son indicates the relationship is total on Man (every Man participates) and used for identification
 - His Name. His Name is now a **discriminant**; this is indicated by **double underline**



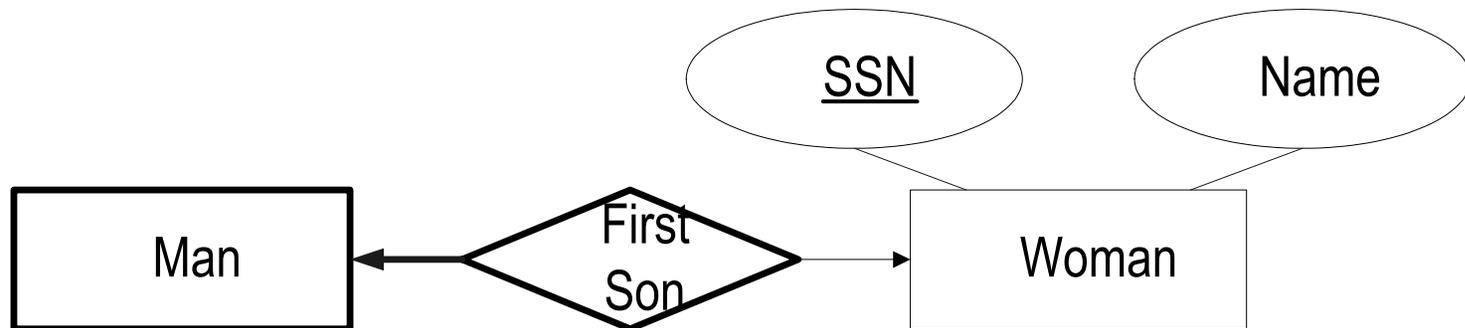
Man As A Weak Entity

- ◆ We need to specify for a weak entity **through which relationship it is identified; this done by using thick lines**
- ◆ Otherwise we do not know whether Man is identified through Son or through Works



Man As A Weak Entity

- ◆ Sometimes a discriminant is not needed
- ◆ We are only interested in men who happen to be first sons of women
- ◆ Every Woman has at most one First Son
- ◆ So we do not need to have Name for Man (if we do not want to store it, but if we do store it, it is not a discriminant)



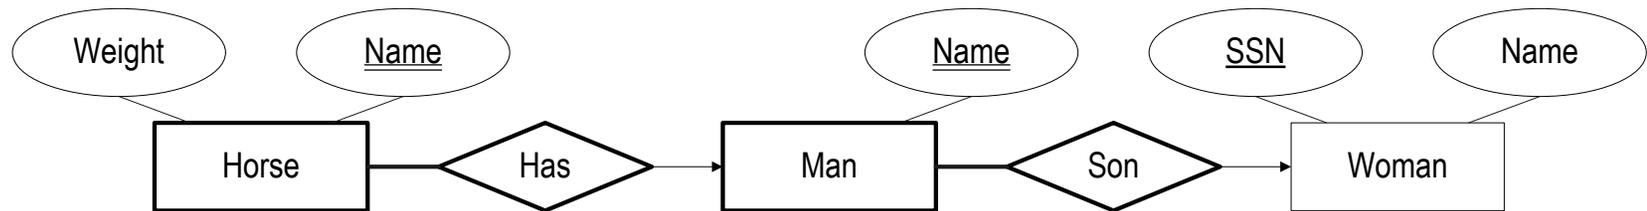
- ◆ Note an arrow to the left: each woman has at most one first son

Man As A Weak Entity

- ◆ In general, more than one attribute may be needed as a discriminant
- ◆ For example, let us say that man has both first name and middle name
- ◆ A mother may give two sons the same first name or the same middle name
- ◆ A mother will never give two sons the same first name and the same middle name
- ◆ The pair (first name, middle name) together form a discriminant

From Weaker To Stronger

- ◆ There can be several levels of “weakness”
- ◆ Here we can say that a horse named “Speedy” belongs to Bob, whose mother is a woman with SSN 072-45-9867



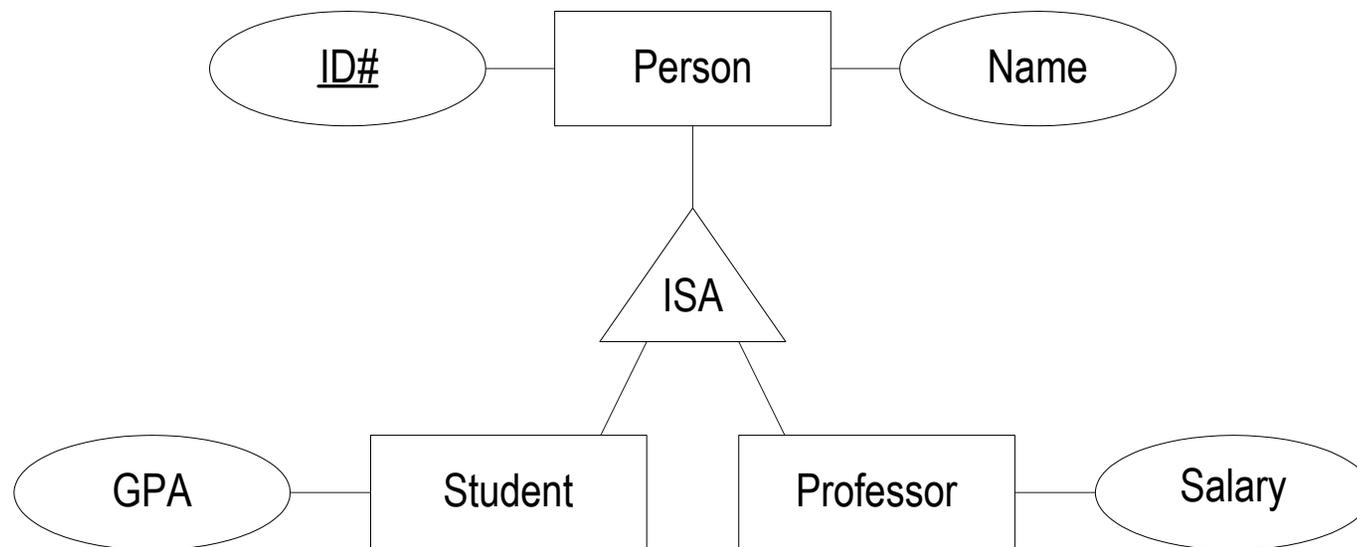
- ◆ A woman can have several sons, each of whom can have several horses

The ISA Relationship

- ◆ For certain purposes, we consider subsets of an entity set
- ◆ The subset relationship between the set and its subset is called **ISA**, meaning “is a”
- ◆ Elements of the subset, of course, have all the attributes and relationships as the elements of the set: they are in the “original” entity set
- ◆ In addition, they may participate in relationships and have attributes that make sense for them
 - But do not make sense for every entity in the “original” entity set
- ◆ ISA is indicated by a triangle
- ◆ The elements of the subset are weak entities, as we will note next

The ISA Relationship

- ◆ Example: A subset that has an attribute that the original set does not have
- ◆ We look at all the persons associated with a university
- ◆ Some of the persons happen to be professors and some of the persons happen to be students



The ISA Relationship

- ◆ Professor is a weak entity because it cannot be identified by its own attributes (here: Salary)
- ◆ Student is a weak entity because it cannot be identified by its own attributes (here: GPA)
- ◆ They do not have discriminants, nothing is needed to identify them in addition to the primary key of the strong entity (Person)

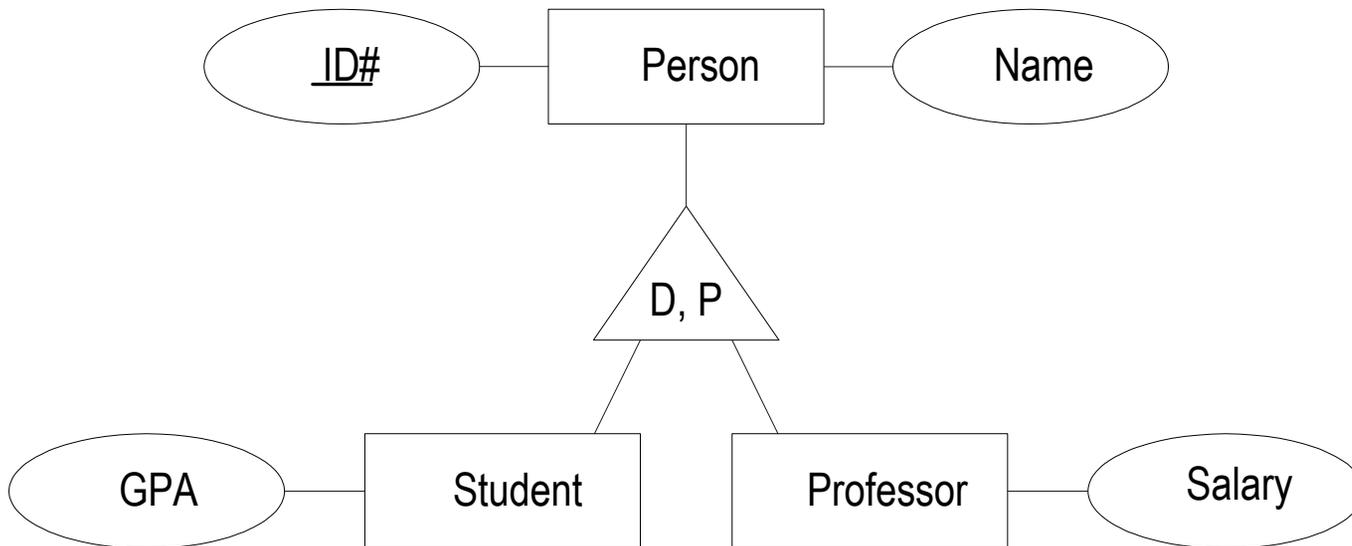
- ◆ The set and the subsets are sometimes referred to as **class** and **subclasses**

The ISA Relationship

- ◆ A person associated with the university (and therefore in our database) can be in general
 - Only a professor
 - Only a student
 - Both a professor and a student
 - Neither a professor nor a student
- ◆ A specific ISA could be
 - **Disjoint**: no entity could be in more than one subclass
 - **Overlapping**: an entity could be in more than one subclass
 - **Total**: every entity has to be in at least one subclass
 - **Partial**: an entity does not have to be in any subclass
- ◆ This could be specified by replacing “ISA” in the diagram by an appropriate letter
- ◆ If nothing stated, then no restriction, so effectively O,P

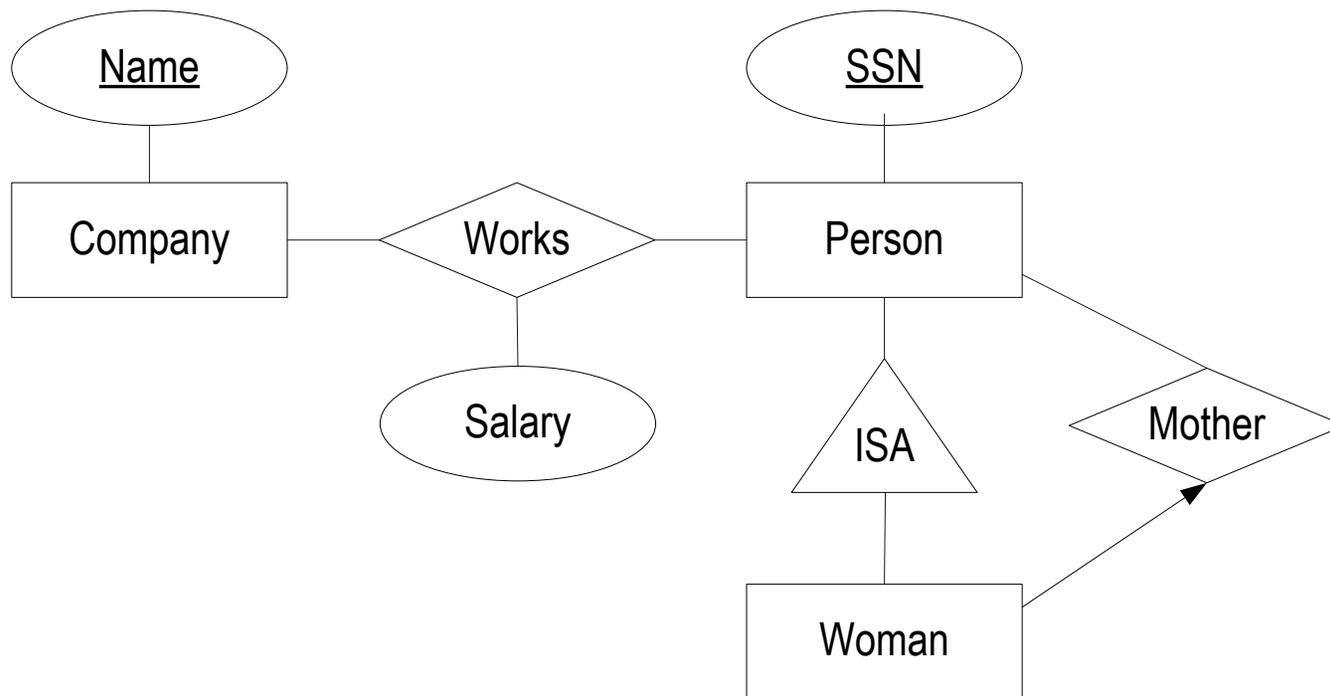
The ISA Relationship

- ◆ Some persons are professors
- ◆ Some persons are students
- ◆ Some persons are neither professors nor students
- ◆ No person can be both a professor and a student



The ISA Relationship

- ◆ Example: subsets participating in relationships modeling the assumed semantics more clearly (every person has one woman who is the birth mother)



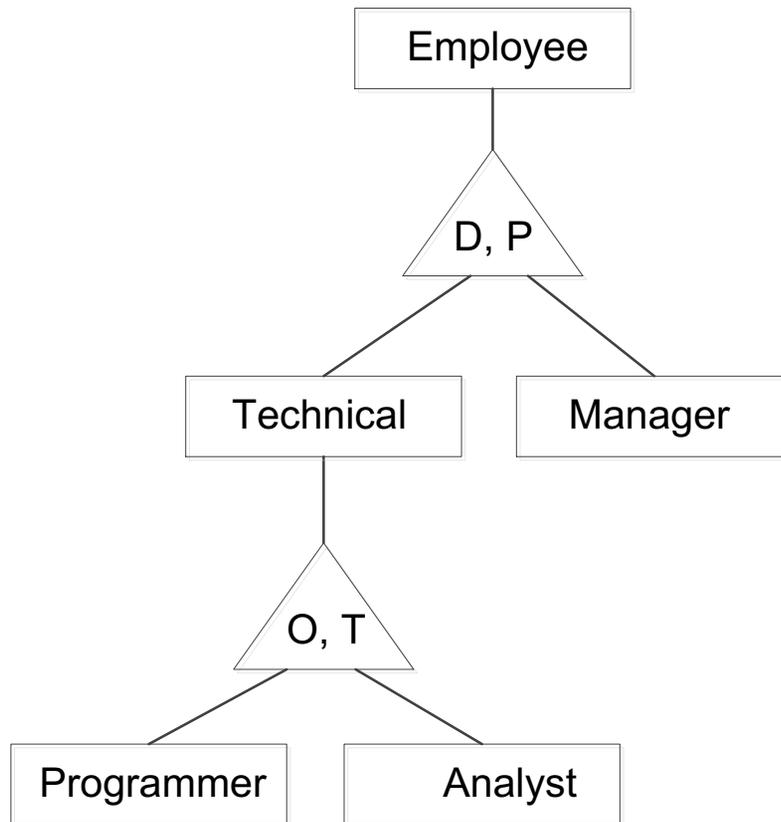
The ISA Relationship

- ◆ ISA is really a superclass/subclass relationship
- ◆ ISA could be **specialization**: subsets are made out of the “more basic” set
- ◆ ISA could be **generalization**: a superset is made of “more basic” sets
- ◆ Again, the diagram could be annotated to indicate this

A More Complex Example

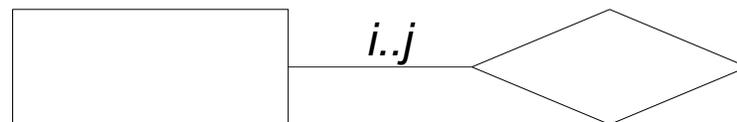
- ◆ We have several types of employees
 - Managers
 - Programmers
 - Analysts
 - Other
- ◆ An employee can be one of the following
 - Manager
 - Programmer and/or Analyst
 - Other
- ◆ The 3 sets are disjoint, that is
 - Manager cannot be Programmer or Analyst, or Other
 - Other cannot be Manager, Programmer, or Analyst
- ◆ All Employees have some share properties
- ◆ It is convenient to group Programmers and Analysts together as they have some shared properties

A Sketch of an ER Diagram



Cardinality Constraints

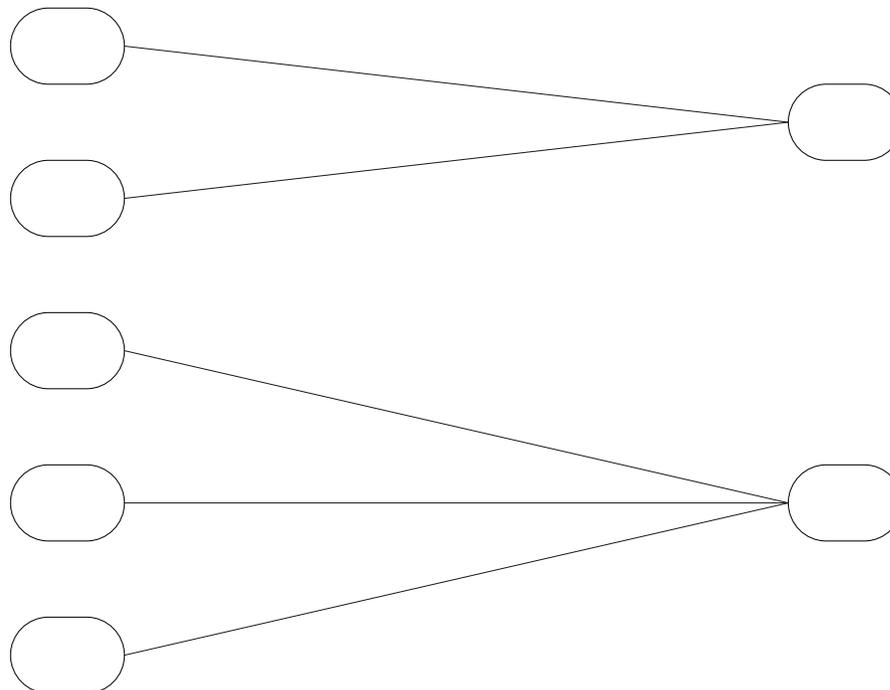
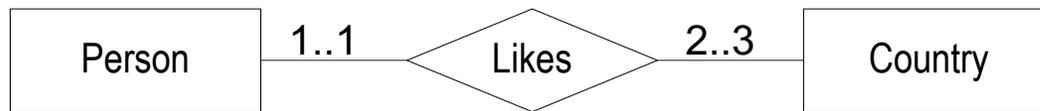
- ◆ We can specify how many times each entity from some entity set can participate in some relationship, in every instance of the database
- ◆ In general we can say that
 - This number is in the interval $[i, j]$, $0 \leq i \leq j$, with i and j integers, denoted by $i..j$; or
 - This number is at the interval $[i, \infty)$, denoted by $i..*$
- ◆ $0..*$ means no constraint
- ◆ No constraint can also be indicated by not writing out anything



- ◆ Note the specific convention we will be using, some people use other conventions for cardinality constraints

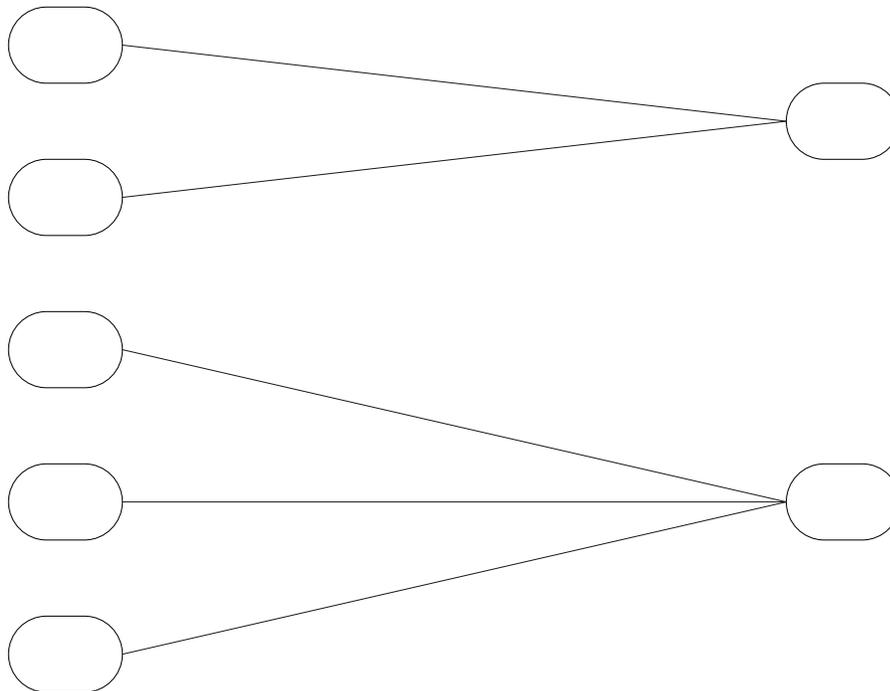
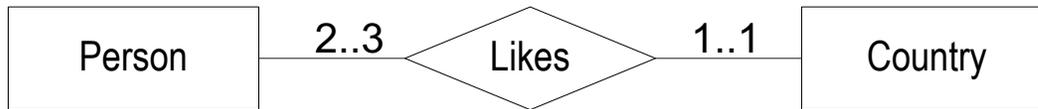
Cardinality Constraints

- ◆ Every person likes exactly 1 country
- ◆ Every country is liked by 2 or 3 persons



Note on Cardinality Constraints

- ◆ Every person likes exactly 1 country
- ◆ Every country is liked by 2 or 3 persons
- ◆ **Sometimes (but not by us) the opposite convention is used**

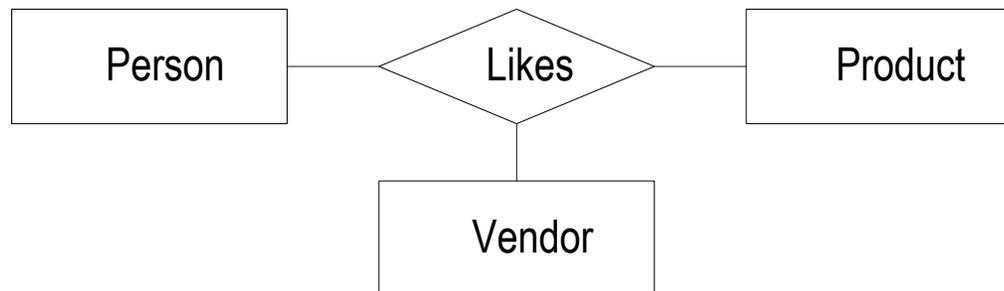


Cardinality Constraints

- ◆ Returning to an old example without specifying which entities actually exist



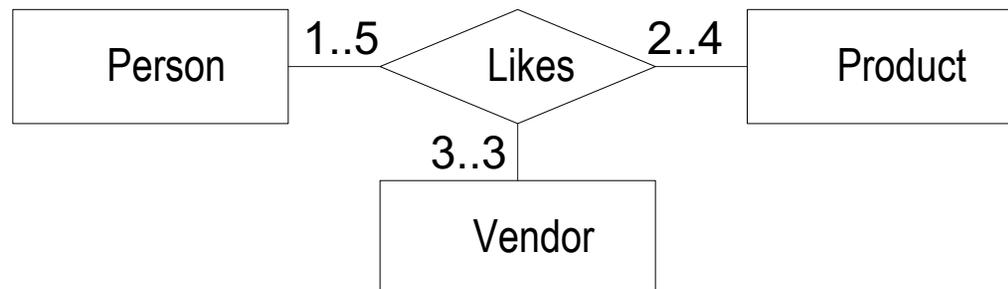
- ◆ We have a relationship: Likes
- ◆ A typical “participation” in a relationship would be that Chee, IBM, Computer participate in it



Cardinality Constraints



- ◆ We want to specify cardinality constraints that **every instance of the database** (that is the **schema**) needs to satisfy
 - Each person participates in between 1 and 5 relationships
 - Each vendor participates in between 3 and 3 (that is exactly 3) relationships
 - Each product participates in between 2 and 4 relationships
- ◆ This is indicated as follows:



Cardinality Constraints

- ◆ A specific instance of the database

Person	Name
	Chee
	Lakshmi
	Marsha

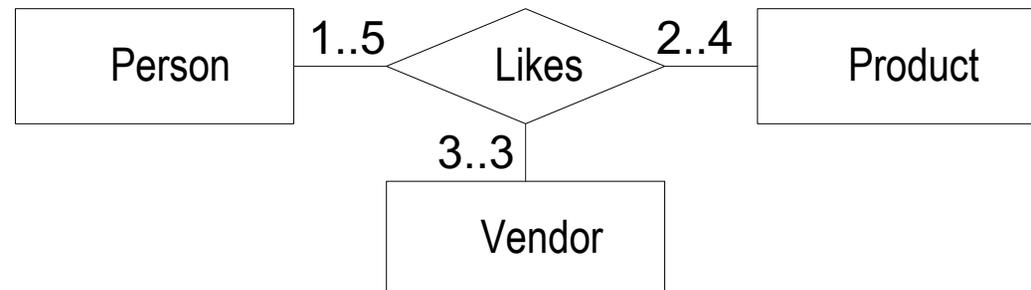
Vendor	Company
	IBM
	Apple

Product	Type
	computer
	monitor

- ◆ If we have the following tuples in the relationship

Chee IBM computer
Lakshmi Apple monitor
Marsha Apple computer
Marsha IBM monitor
Marsha IBM computer
Lakshmi Apple computer

- ◆ Then, it is true that:



Cardinality Constraints

- ◆ Let us confirm that our instance of Likes satisfies the required cardinality constraints
- ◆ Person: required between 1 and 5
 - Chee in 1
 - Lakshmi in 2
 - Marsha in 3
- ◆ Product: required between 2 and 4
 - Monitor in 2
 - Computer in 4
- ◆ Vendor between 3 and 3
 - Apple in 3
 - IBM in 3
- ◆ Note that we do not have to have an entity for every possible permitted cardinality value
 - For example, there is no person participating in 4 or 5 tuples

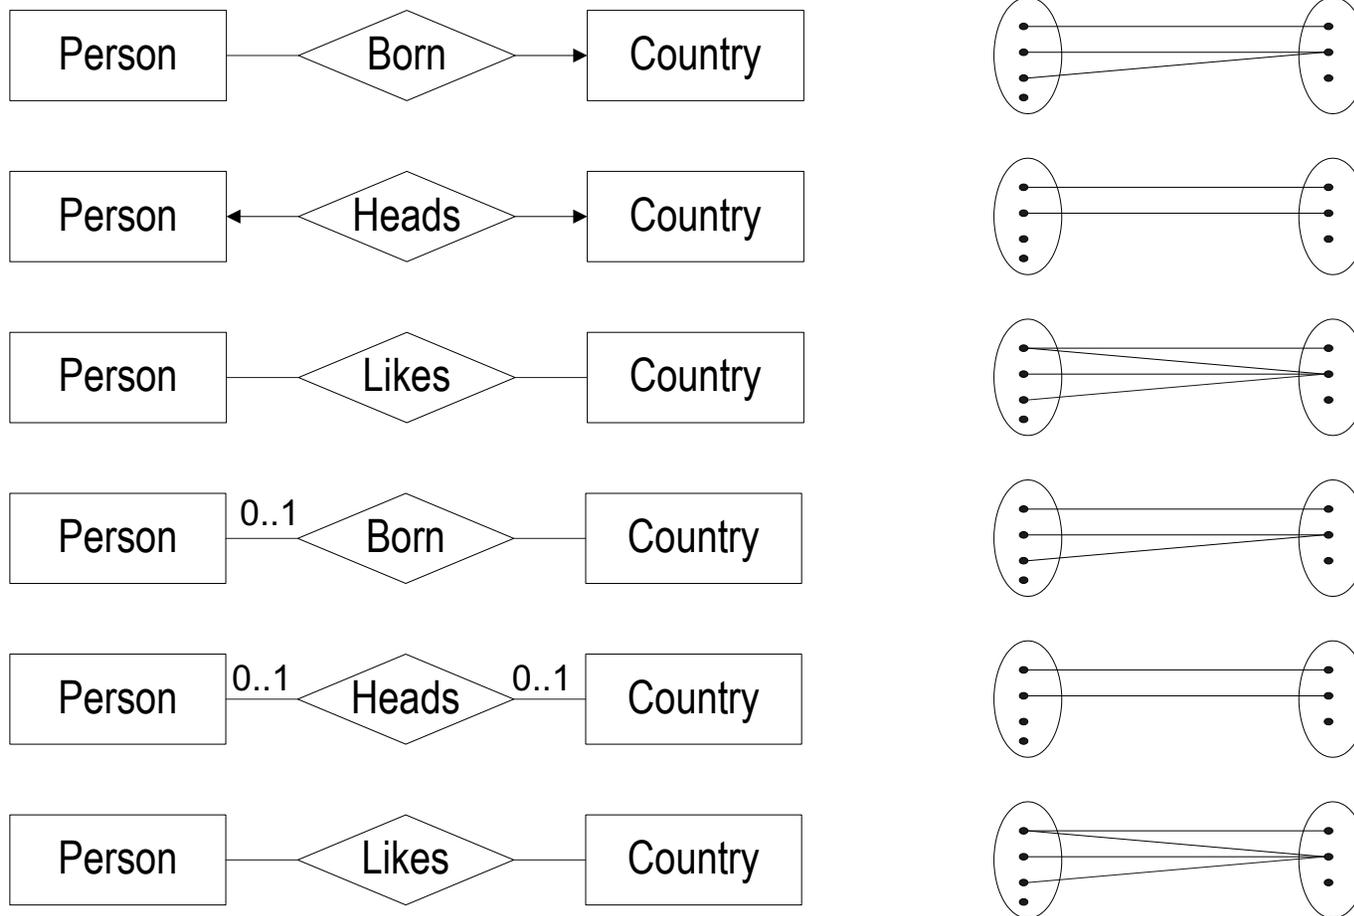
Cardinality Constraints

- ◆ So we can also have, expressing exactly what we had before



Cardinality Constraints

◆ Compare to previous notation



A Case Study

- ◆ Next, we will go through a relatively large example to make sure we know how to use ER diagrams
- ◆ We have a large application to make sure we understand all the points
- ◆ The fragment has been constructed so it exhibits interesting and important capabilities of modeling
- ◆ It will also review the concepts we have studied earlier
- ◆ It is chosen based on its suitability to practice modeling using the power of ER diagrams
- ◆ It will also exercise various points, to be discussed later on how to design good relational databases

Our Application

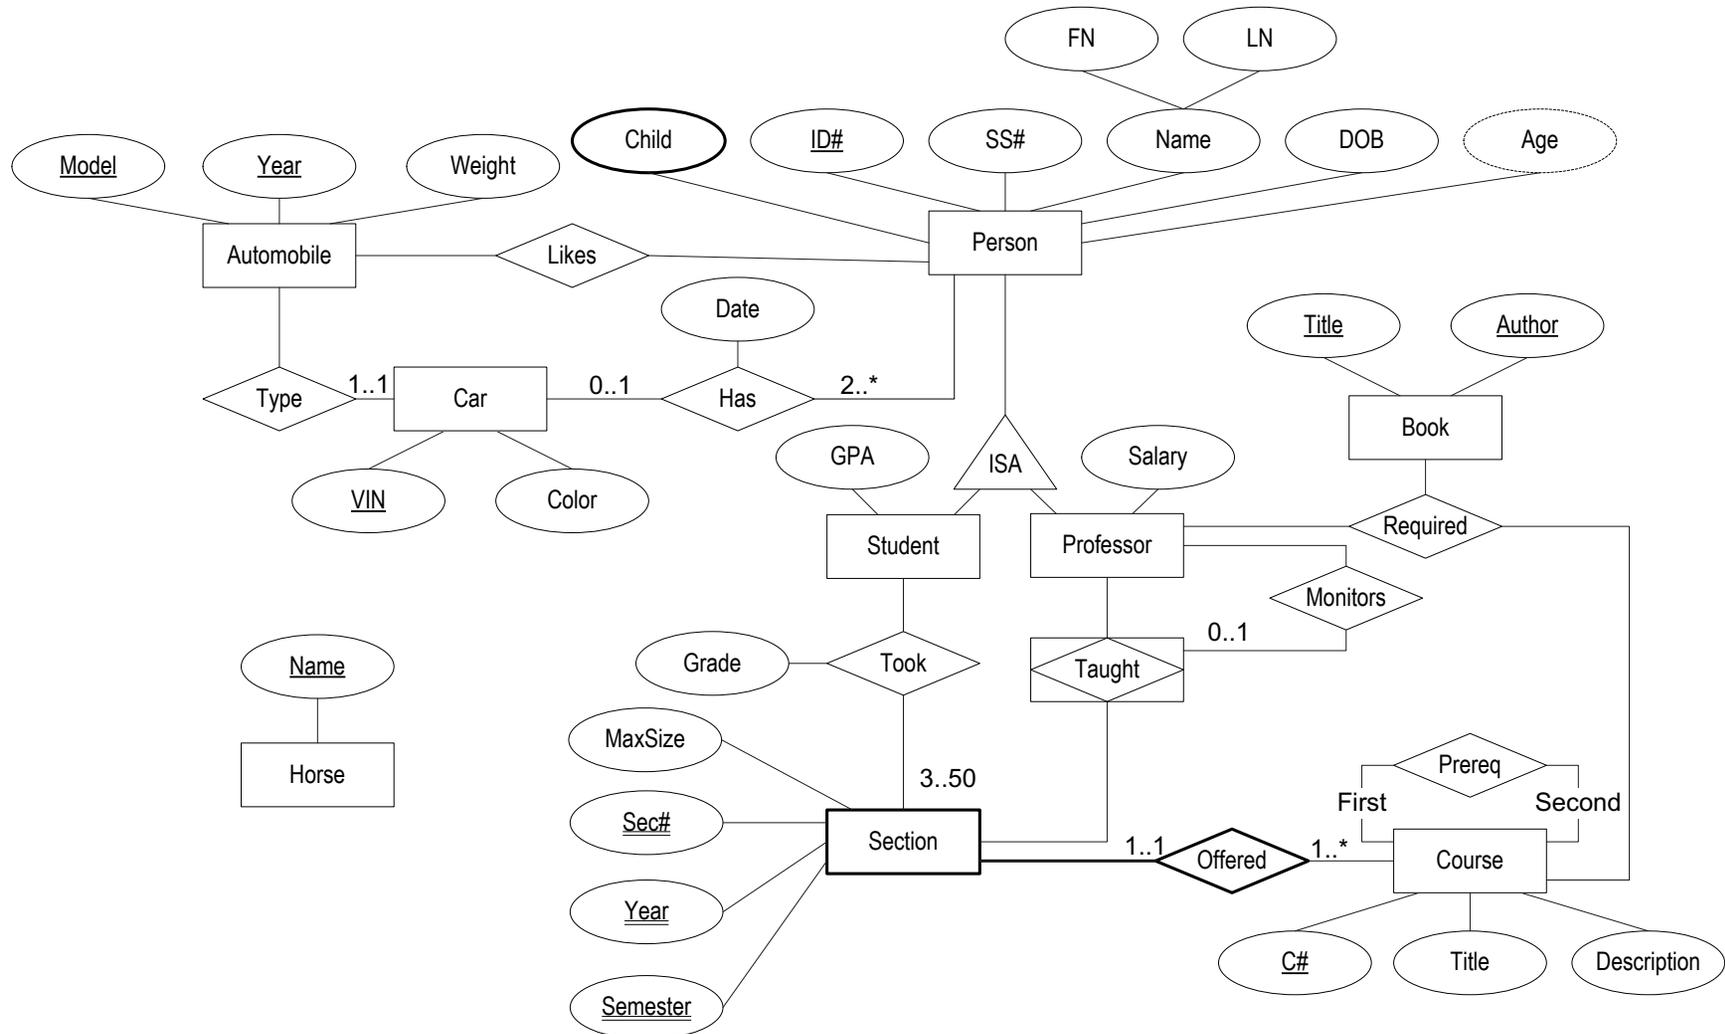
- ◆ We are supposed to design a database for a university
- ◆ We will look at a small fragment of the application and will model it as an entity relationship diagram annotated with comments, as needed to express additional features
- ◆ But it is still a reasonable “small” database
- ◆ In fact, much larger than what is commonly discussed in a course, but more realistic for modeling real applications

Our Application

- ◆ Our understanding of the application will be described in a narrative form
- ◆ While we do this, we construct the ER diagram
- ◆ For ease of exposition (technical reasons only: limitations of the projection equipment) we look at the resulting ER diagram and construct it in pieces
- ◆ We will pick some syntax for annotations, as this is not standard
- ◆ One may try and write the annotations on the diagram itself using appropriate phrasing, but this will make our example too cluttered

Building The ER Diagram

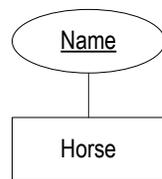
- ◆ We describe the application in stages, getting:



Horse

- ◆ *Horse*; entity set
- ◆ Attributes:
 - *Name*
- ◆ Constraints
 - Primary Key: Name

Our ER Diagram



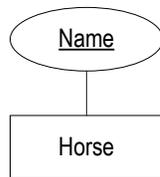
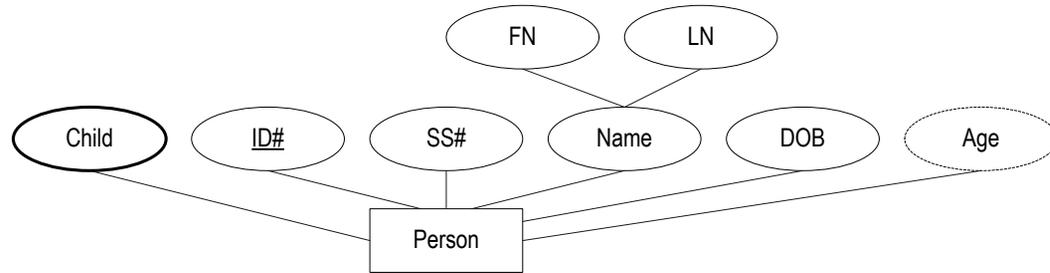
Horse

- ◆ We should specify what is the domain of each attribute, in this case, Name only
- ◆ We will generally not do it in our example, as there is nothing interesting in it
 - We could say that Name is an alphabetic string of at most 100 characters, for example

Person

- ◆ **Person**; entity set
- ◆ Attributes:
 - **Child**; a multivalued attribute
 - **ID#**
 - **SS#**
 - **Name**; composite attribute, consisting of
 - **FN**
 - **LN**
 - **DOB**
 - **Age**; derived attribute (we should state how it is computed)
- ◆ Constraints
 - Primary Key: ID#
 - Unique: SS# (Note that this must be stated in words as we do not have a way of marking the diagram directly)

Our ER Diagram



Person

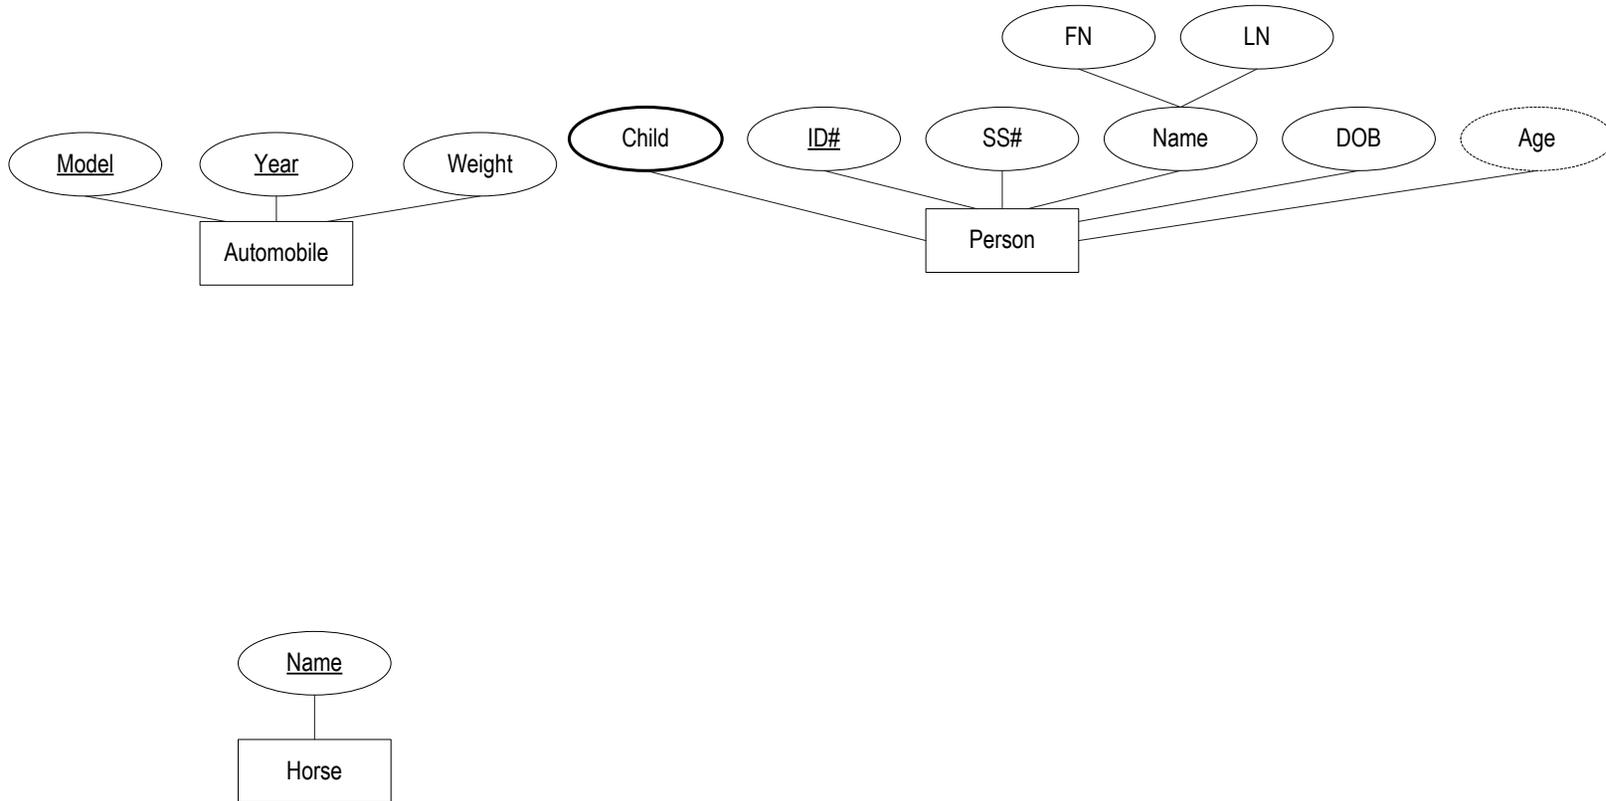
- ◆ Since ID# is the primary key (consisting here of one attribute), we will consistently identify a person using the value of this attribute (for later implementation as a relational database)
- ◆ Since SS# is unique, no two persons will have the same SS# (and we need to tell the database that property, so it can be enforced)

Automobile

- ◆ ***Automobile***; entity set
- ◆ Attributes:
 - ***Model***
 - ***Year***
 - ***Weight***
- ◆ Constraints
 - Primary Key: Model, Year

- ◆ Note: Automobile is a “catalog entry”
 - It is not a specific “physical car”

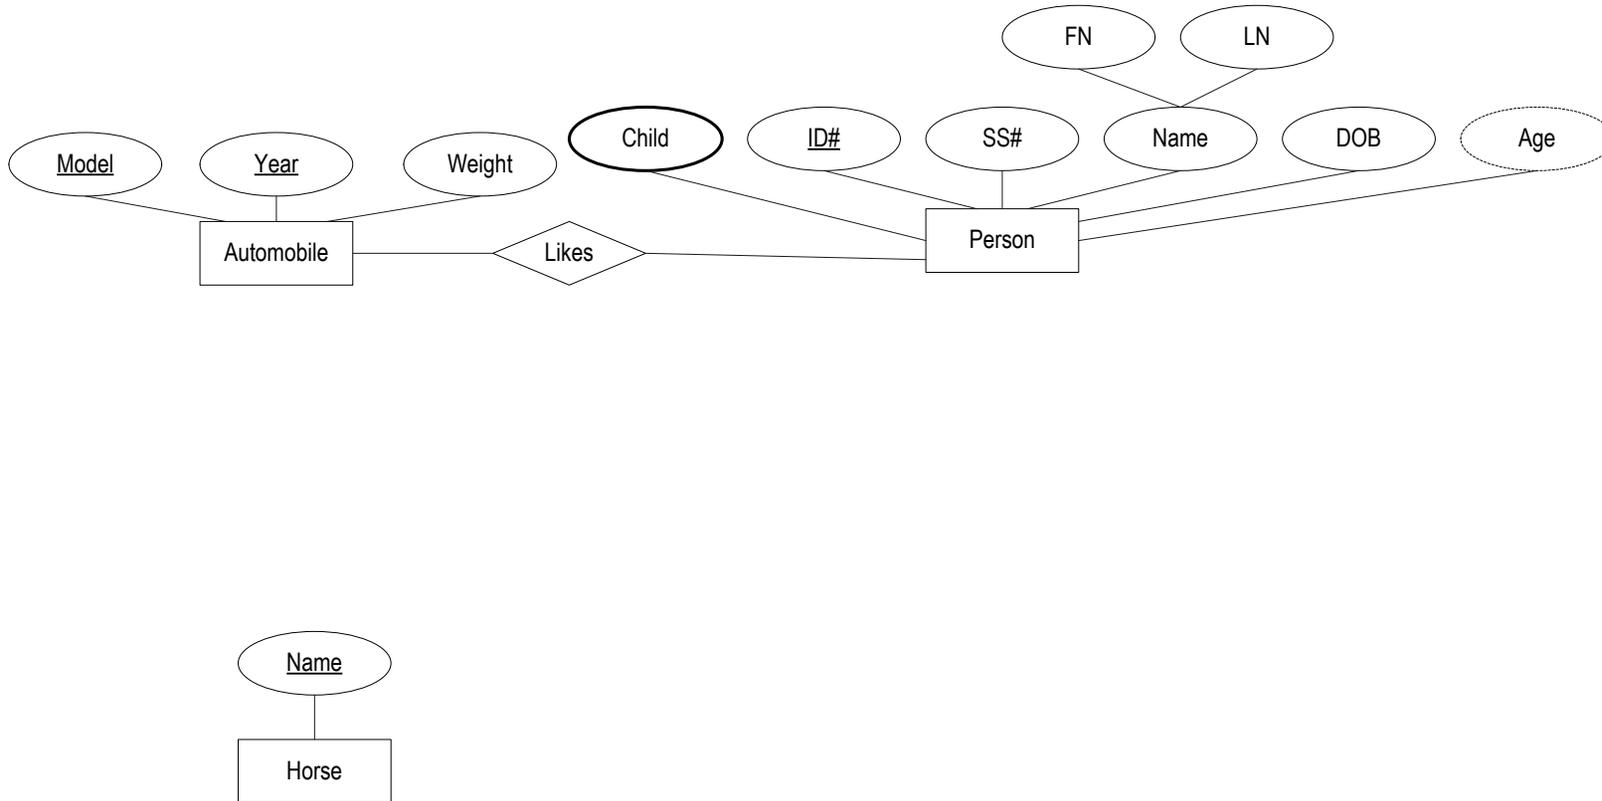
Our ER Diagram



Likes

- ◆ **Likes**; relationship
- ◆ Relationship among/between:
 - **Person**
 - **Automobile**
- ◆ Attributes
- ◆ Constraints

Our ER Diagram



Likes

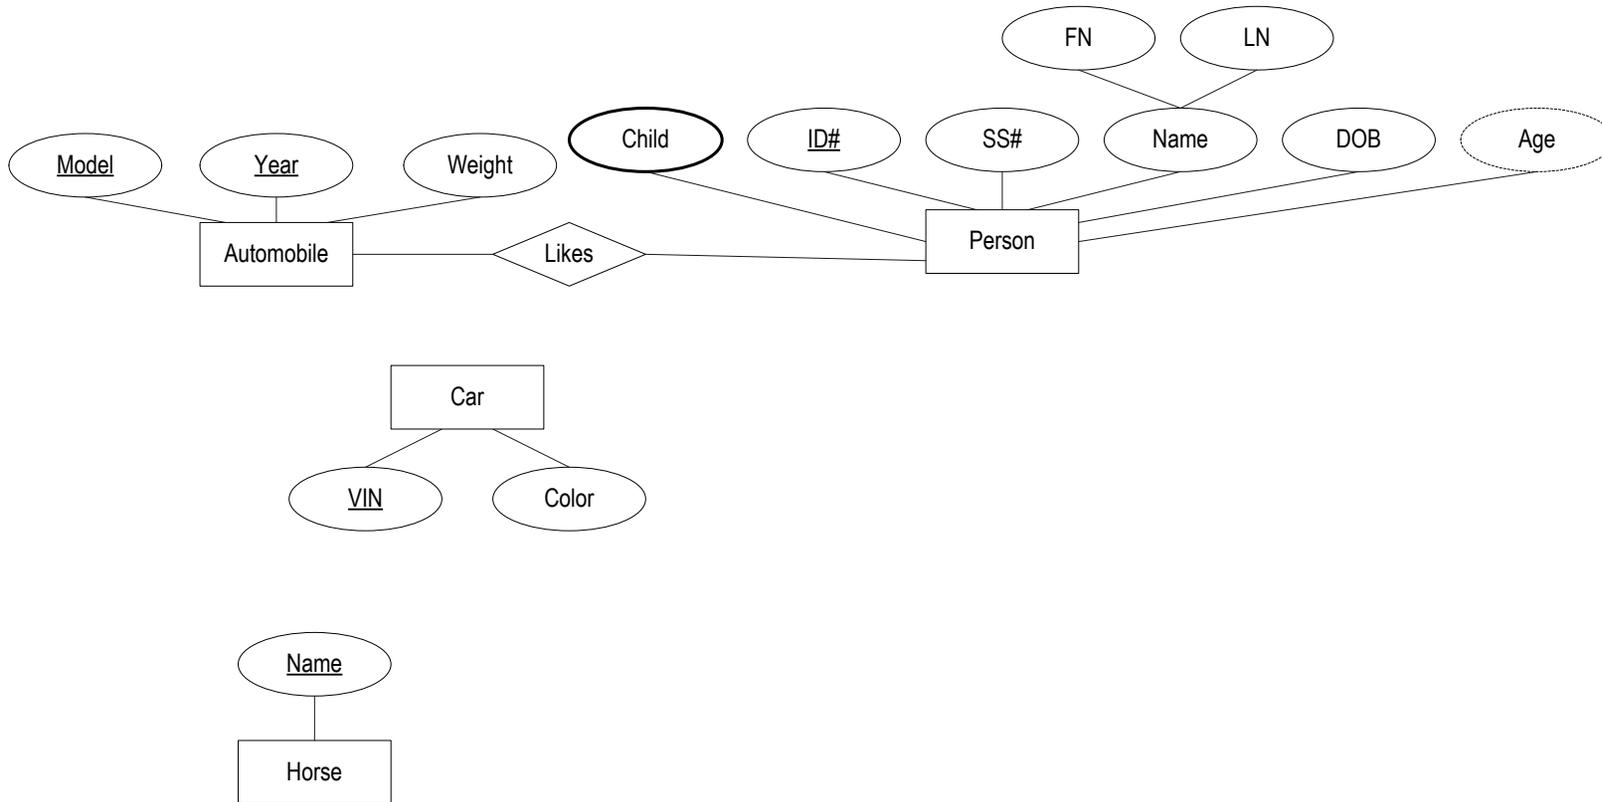
- ◆ This relationship has no attributes
- ◆ This relationship has no constraints
- ◆ This relationship is a general many-to-many relationship (as we have not said otherwise)
- ◆ This relationship does not have any cardinality constraints

Car

- ◆ **Car**; entity set
- ◆ Attributes
 - **VIN**
 - **Color**
- ◆ Constraints
 - Primary Key: VIN

- ◆ Note: Car is a “physical entity”
 - VIN stands for “Vehicle Identification Number,” which is like a Social Security Number for cars

Our ER Diagram

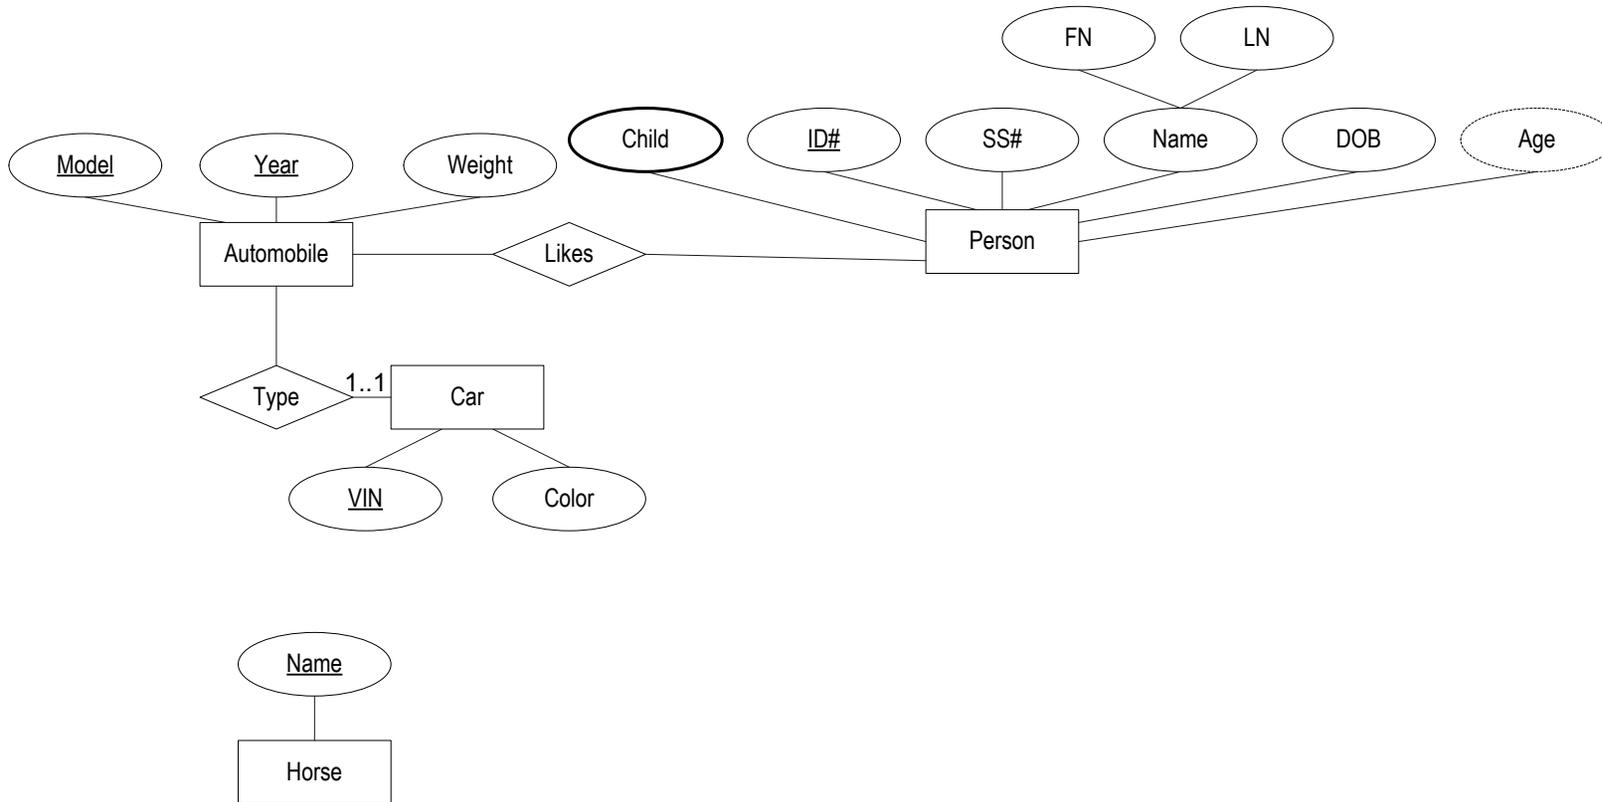


Type

- ◆ **Type**; relationship
- ◆ Relationship among/between:
 - **Automobile**
 - **Car**
- ◆ Attributes
- ◆ Constraints
 - Cardinality: 1..1 between Car and Type

- ◆ This tells us for each physical car what is the automobile catalog entry of which it is an instantiation
 - Each car is an instantiation of a exactly one catalog entry

Our ER Diagram



Type

- ◆ We see that the relationship Type is:
- ◆ Many to one from Car to Automobile
- ◆ It is total not partial

In other words, it is a total function from Car to Automobile

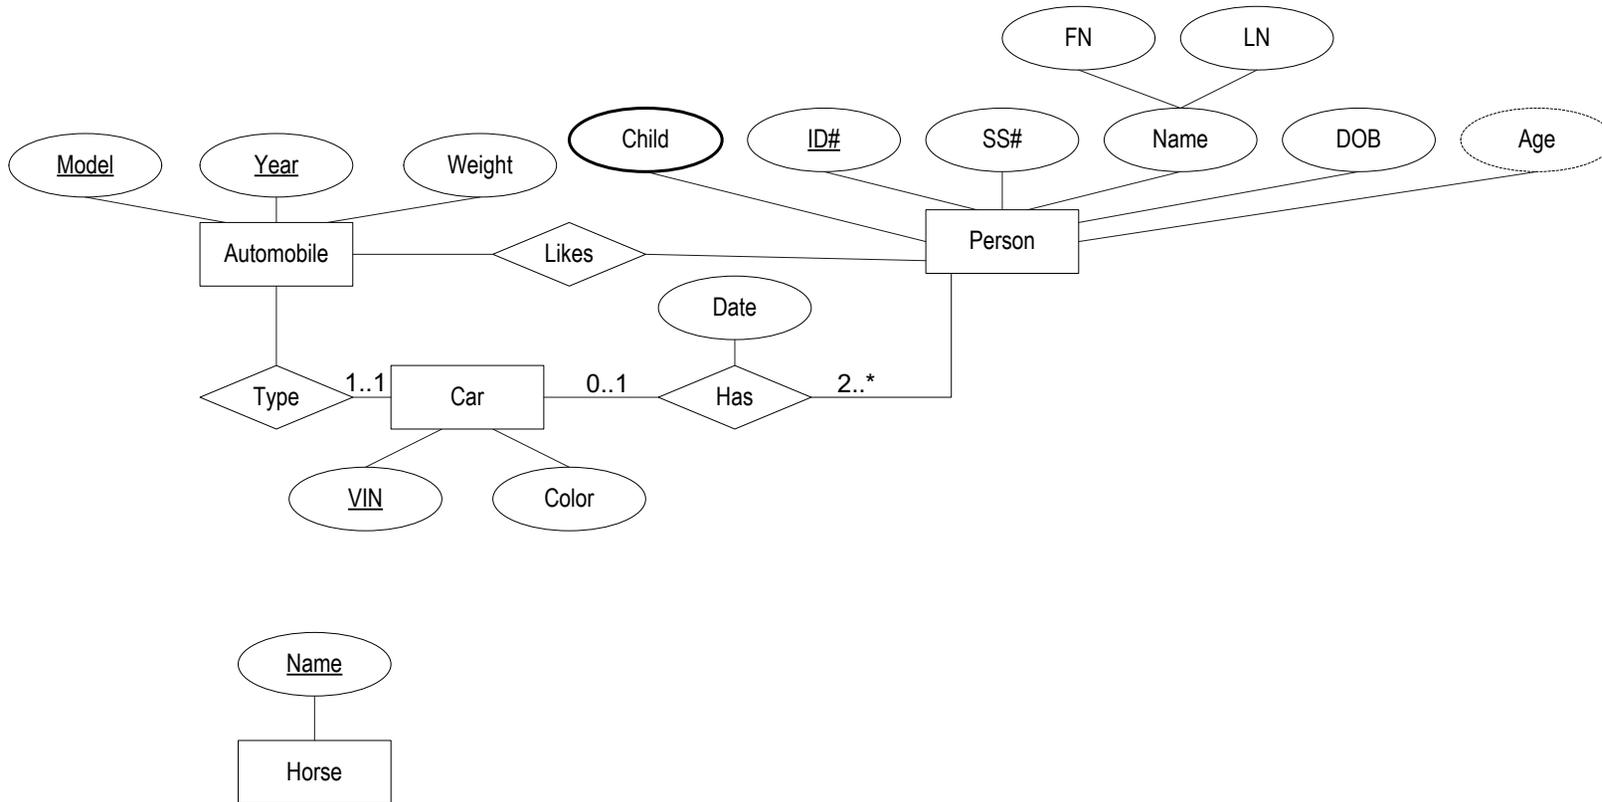
- ◆ Not every Automobile is a “target”
There may be elements in Automobile for which no Car exists

Has

- ◆ **Has**; relationship
- ◆ Relationship among/between
 - **Person**
 - **Car**
- ◆ Attributes
 - **Date**
- ◆ Constraints
 - Cardinality: 2..* between Person and Has
 - Cardinality: 0..1 between Car and Has

- ◆ Date tells us when the person got the car
- ◆ Every person has at least two cars
- ◆ Every car can be had (owned) by at most one person
 - Some cars may have been abandoned

Our ER Diagram



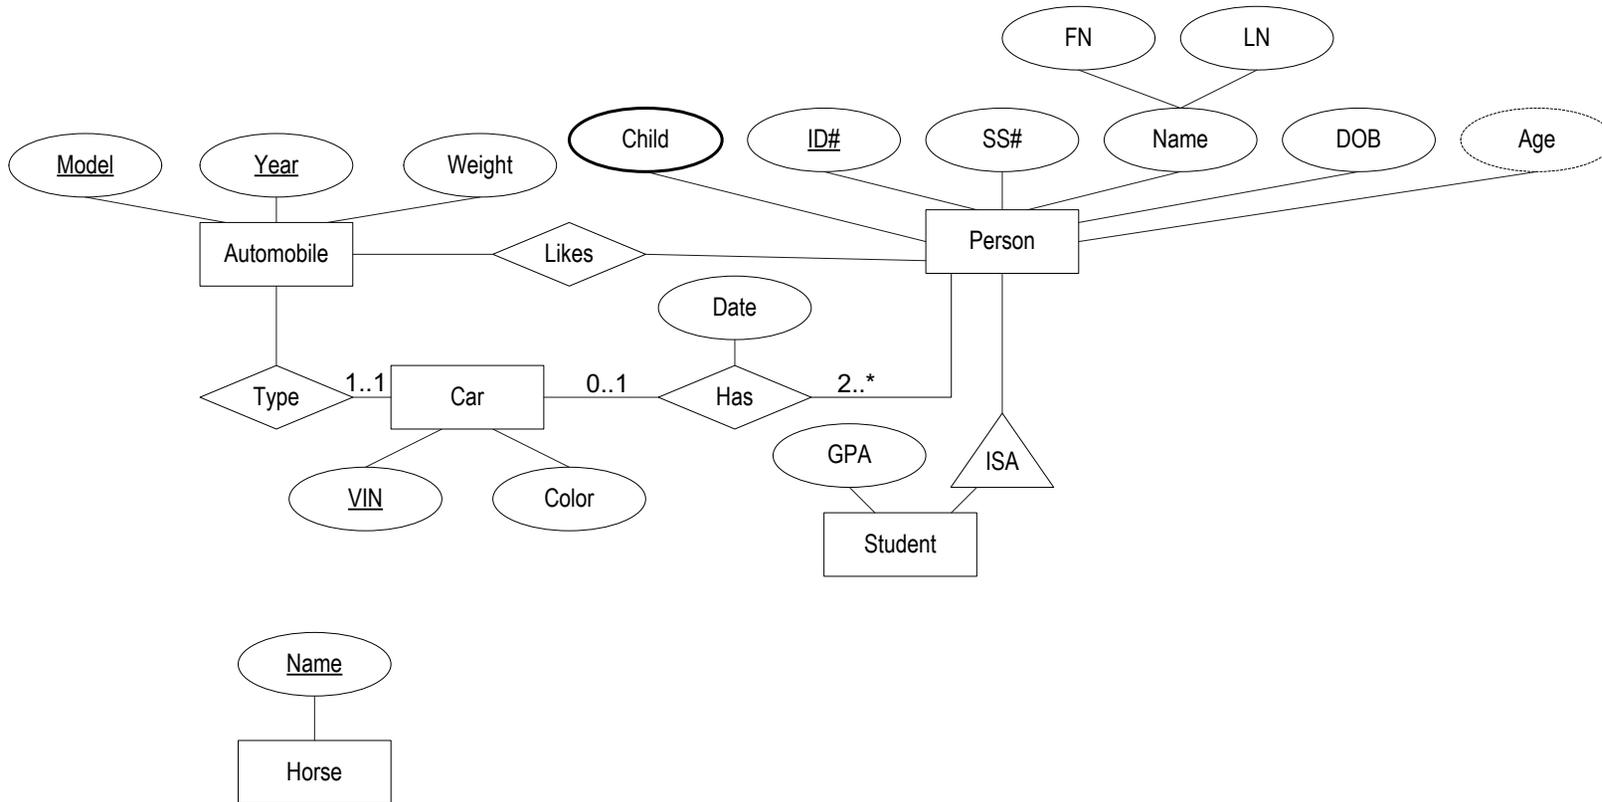
Has

- ◆ We see that Has is a partial function from Car to Person
- ◆ Every Person is a “target” in this function (in fact at least twice)

Student

- ◆ *Student*; entity set
 - ◆ Subclass of Person
 - ◆ Attributes
 - *GPA*
 - ◆ Constraints
-
- ◆ Note that Student is a weak entity
 - It is identified through a person
 - You may think of a student as being an “alias” for some person
“Split personality”

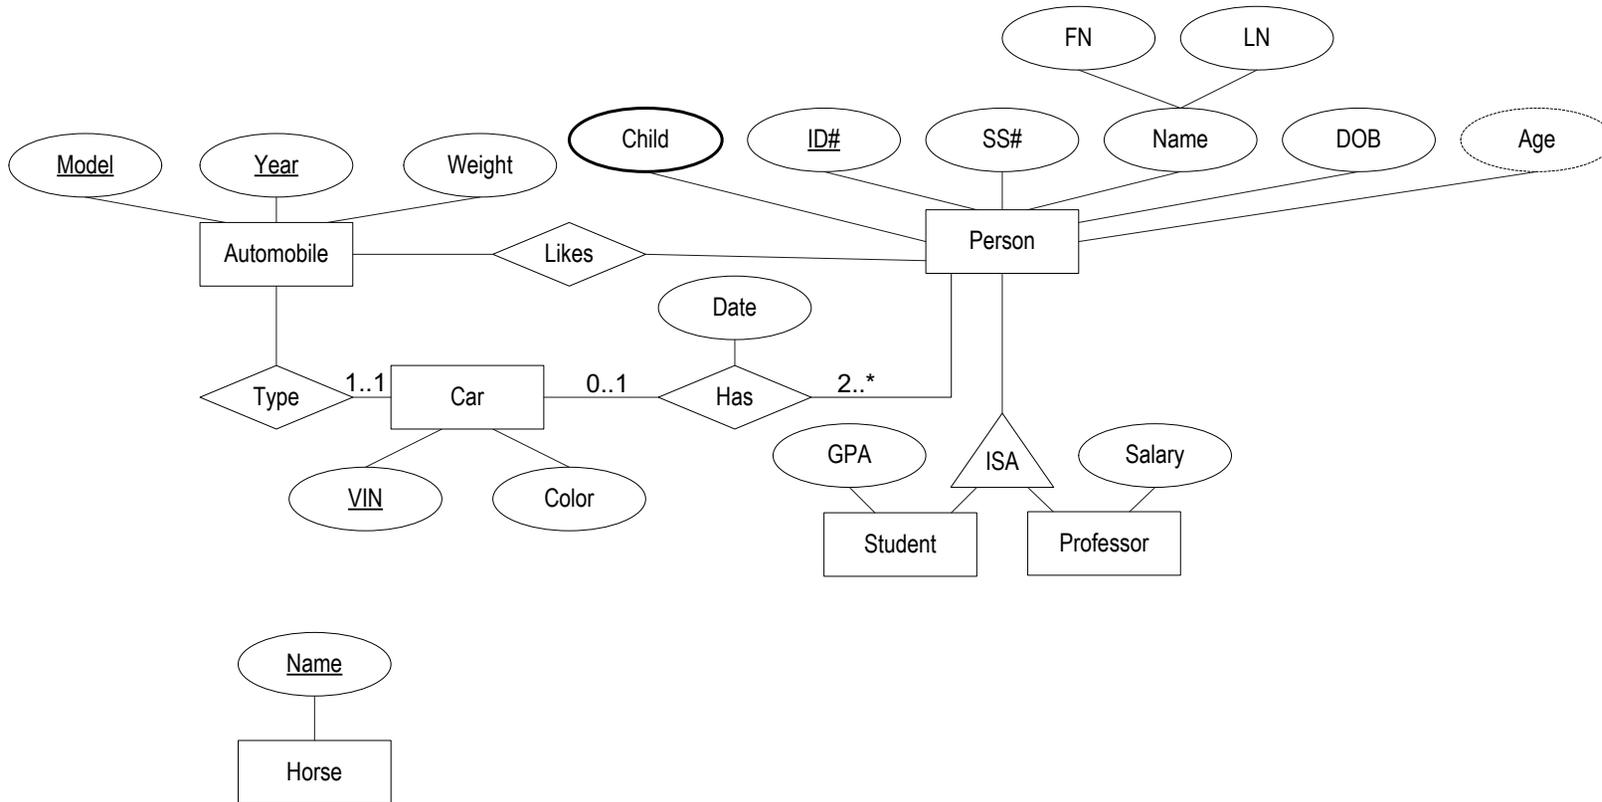
Our ER Diagram



Professor

- ◆ *Professor*; entity set
- ◆ Subclass of Person
- ◆ Attributes
 - *Salary*
- ◆ Constraints

Our ER Diagram

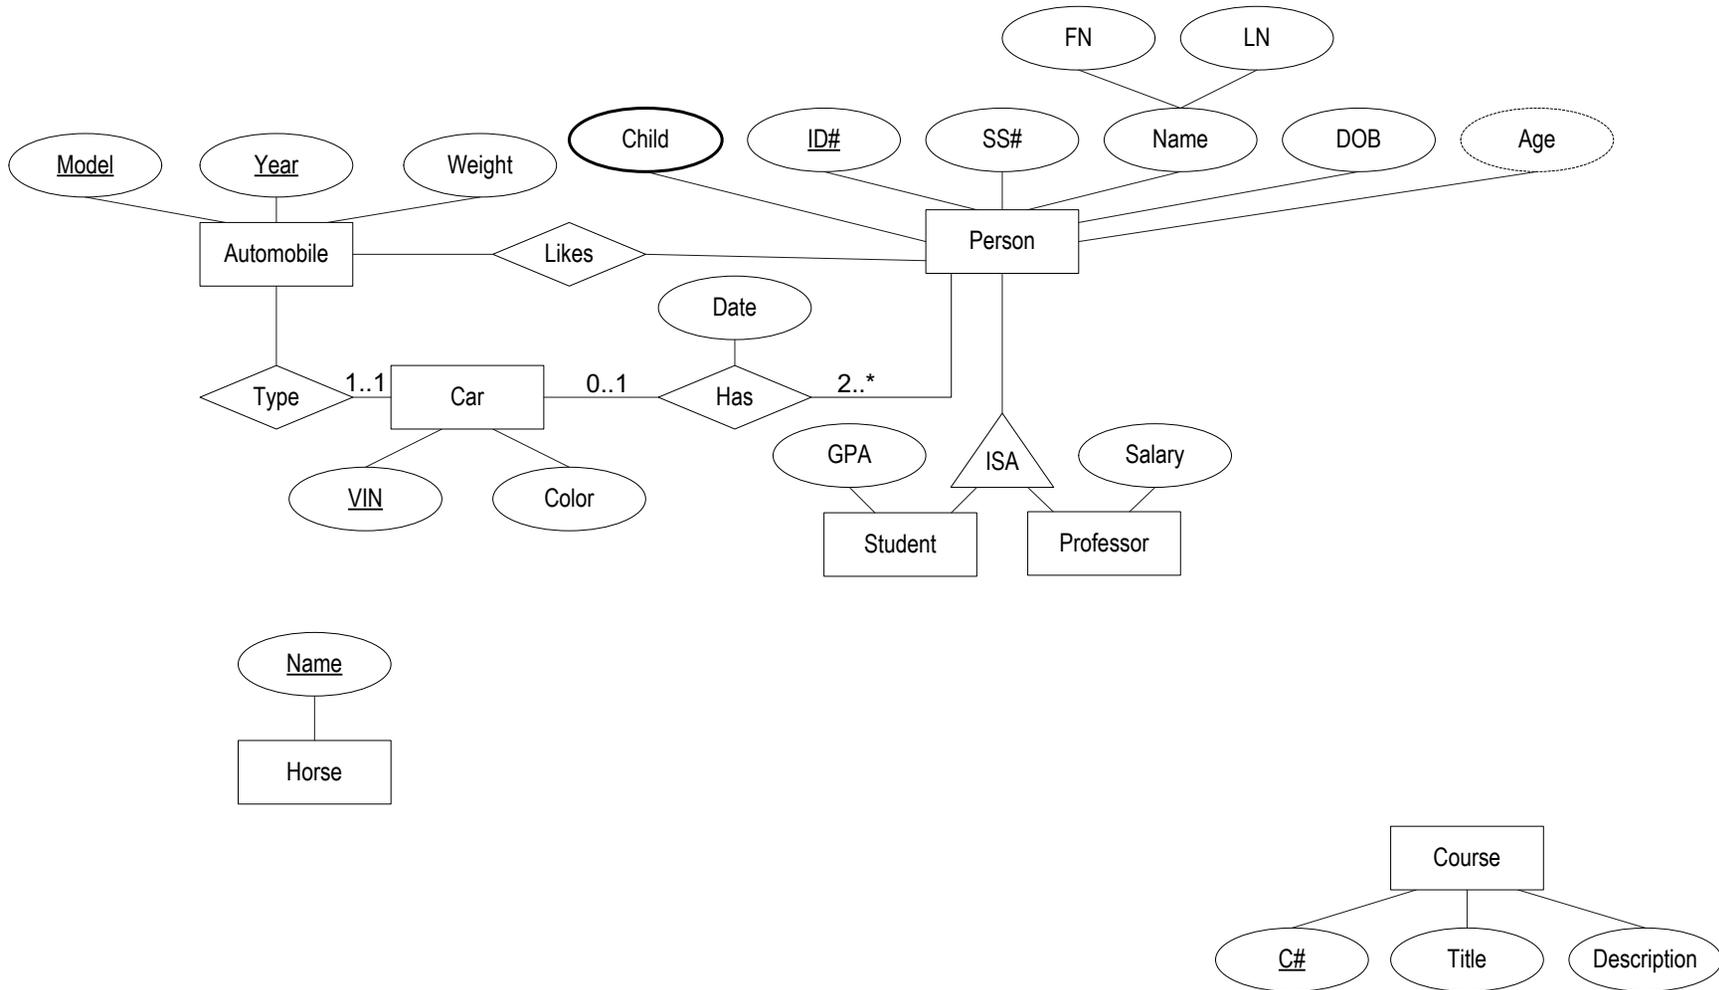


Course

- ◆ **Course**; entity set
- ◆ Attributes:
 - **C#**
 - **Title**
 - **Description**
- ◆ Constraints
 - Primary Key: C#

- ◆ Course is a catalog entry appearing in the bulletin
 - Not a particular offering of a course
 - Example: CSCI-GA.2433 (which is a C#)

Our ER Diagram



Prereq

- ◆ **Prereq**; relationship
- ◆ Relationship among/between:
 - **Course**; role: First
 - **Course**; role: Second
- ◆ Attributes
- ◆ Constraints

- ◆ We have a directed graph on courses, telling us prerequisites for each course, if any
 - To take “second” course every “first” course related to it must have been taken previously
 - We needed the roles first and second, to be clear
 - Note how we model well that prerequisites are not between offerings of a course but catalog entries of courses
 - Note however, that we cannot directly “diagram” that a course cannot be a prerequisite for itself, and similar, so these need to be annotated

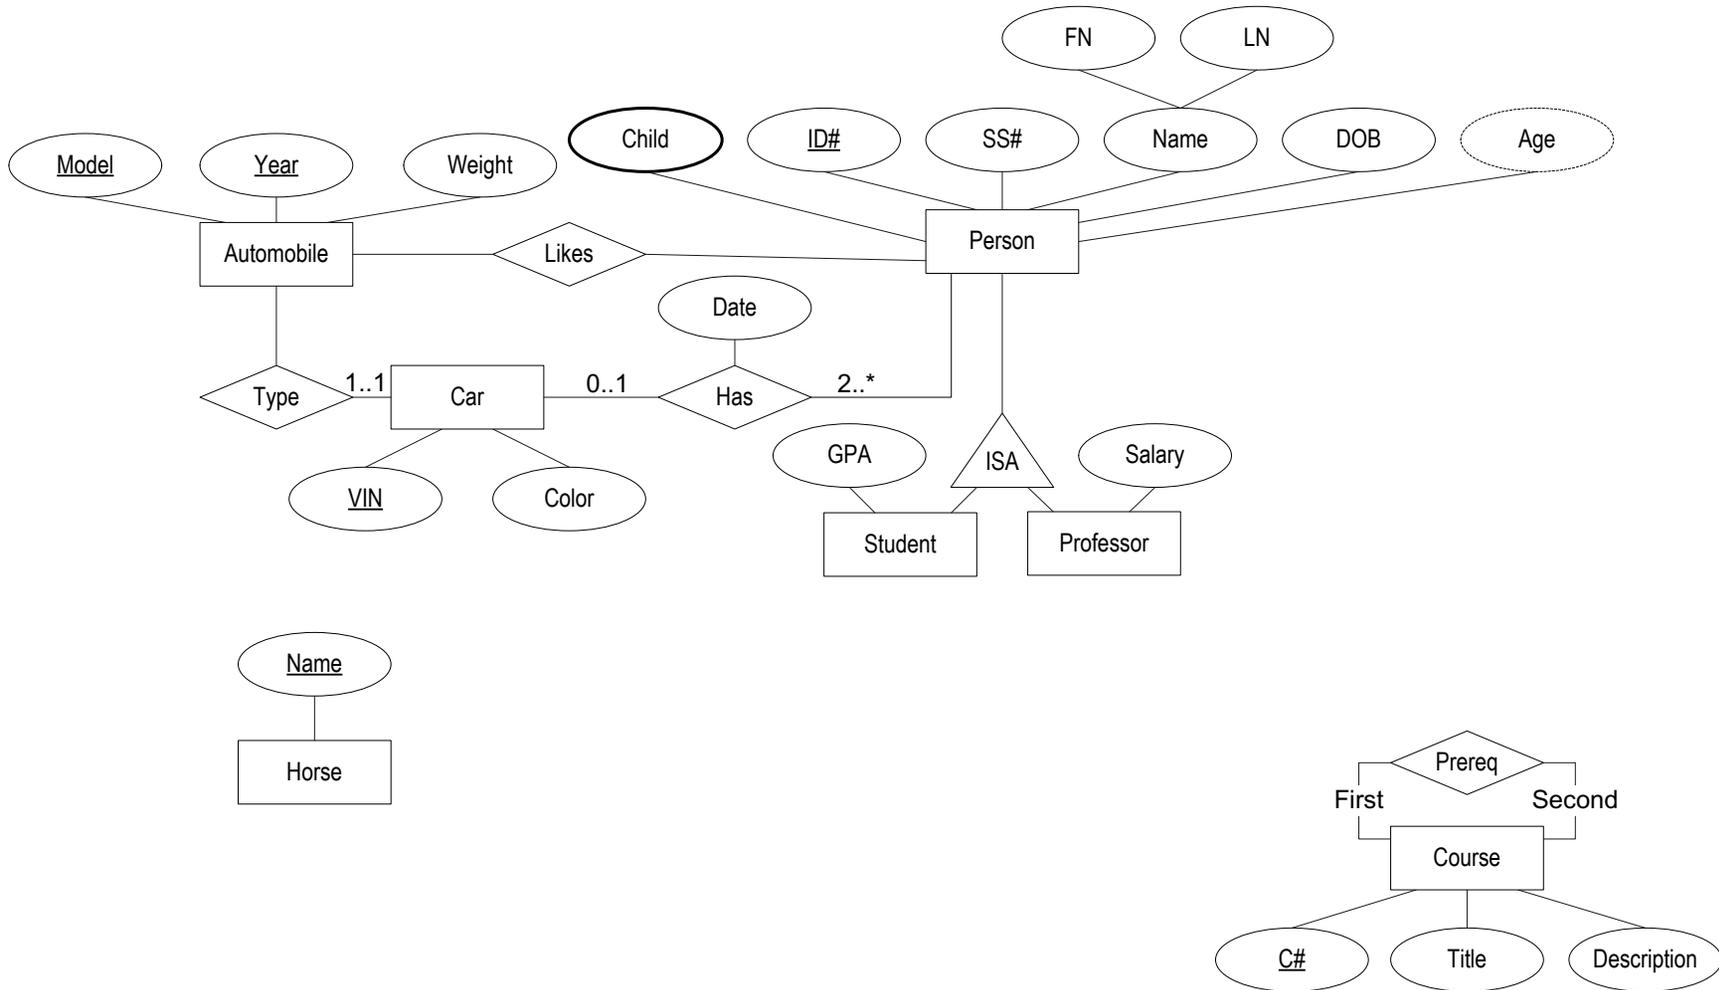
Important Digression

- ◆ Ultimately, we will store (most) relationships as tables
- ◆ So, comparing to our example for Likes, Prereq instance could be

Prereq	First	Second
	101	103
	101	104
	102	104
	104	105
	107	106
	107	108

- ◆ Where we identify the “participating” entities using their primary keys, but renaming them using roles
- ◆ So looking at the table we see that 101 is a prerequisite for 104 but 104 is not a prerequisite for 101

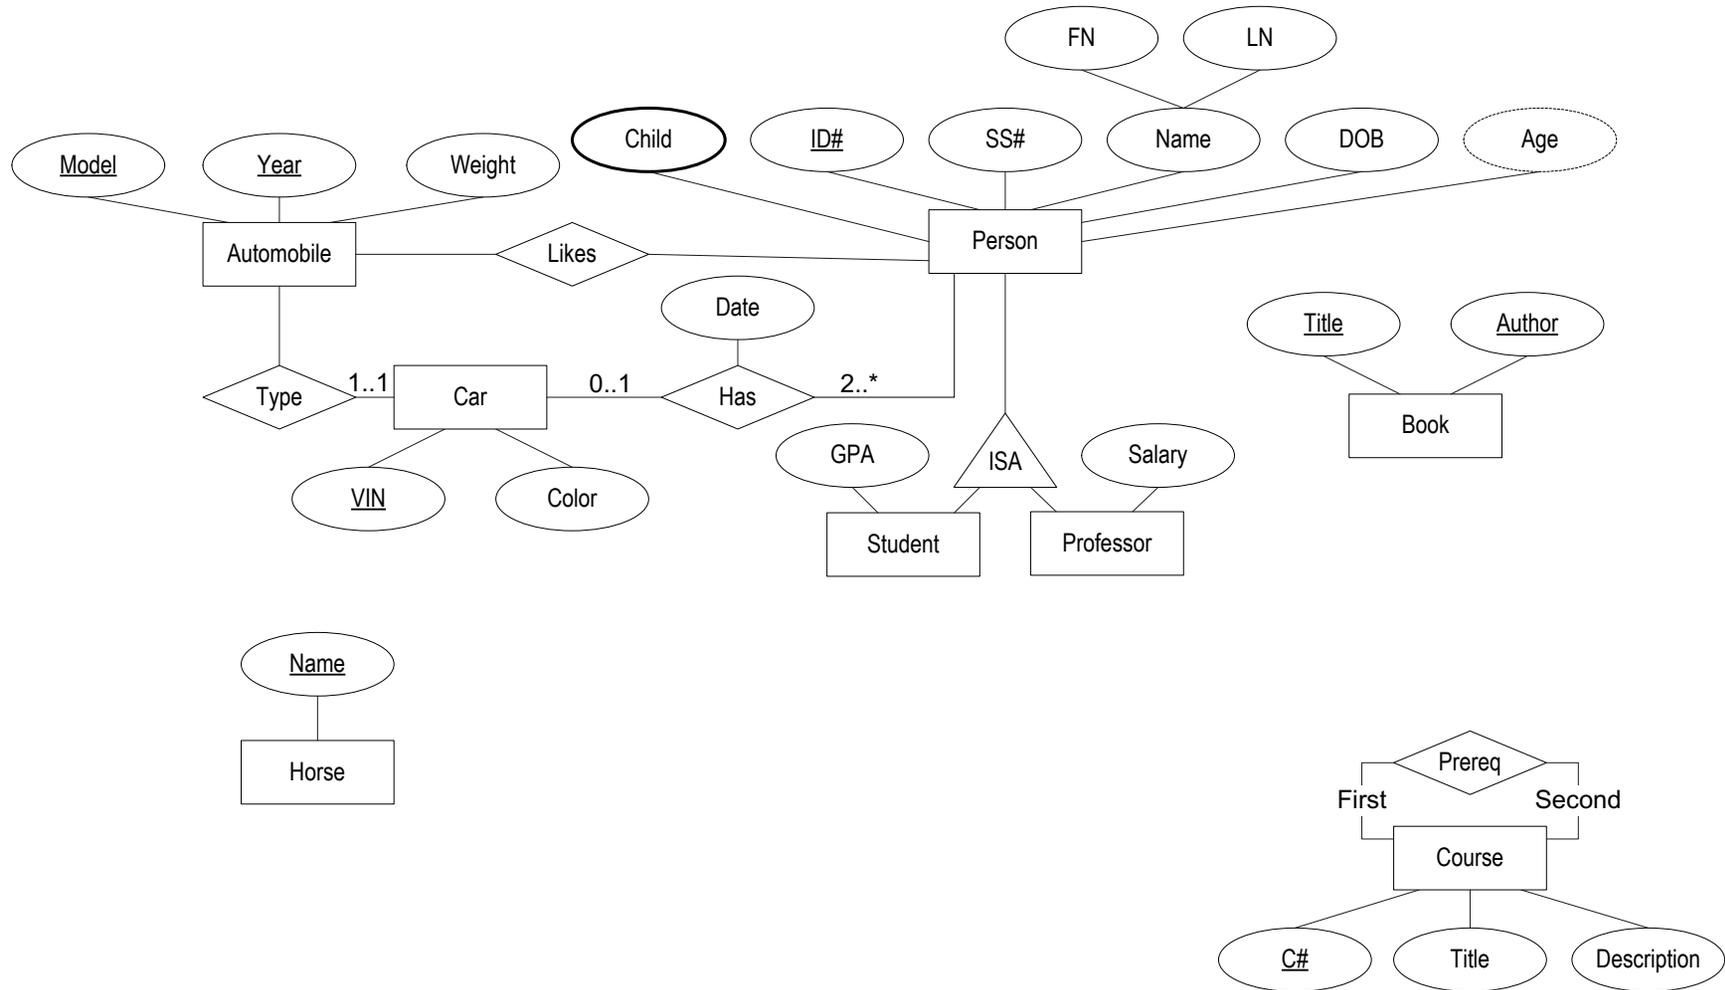
Our ER Diagram



Book

- ◆ ***Book***; entity set
- ◆ Attributes:
 - ***Author***
 - ***Title***
- ◆ Constraints
 - Primary Key: Author, Title

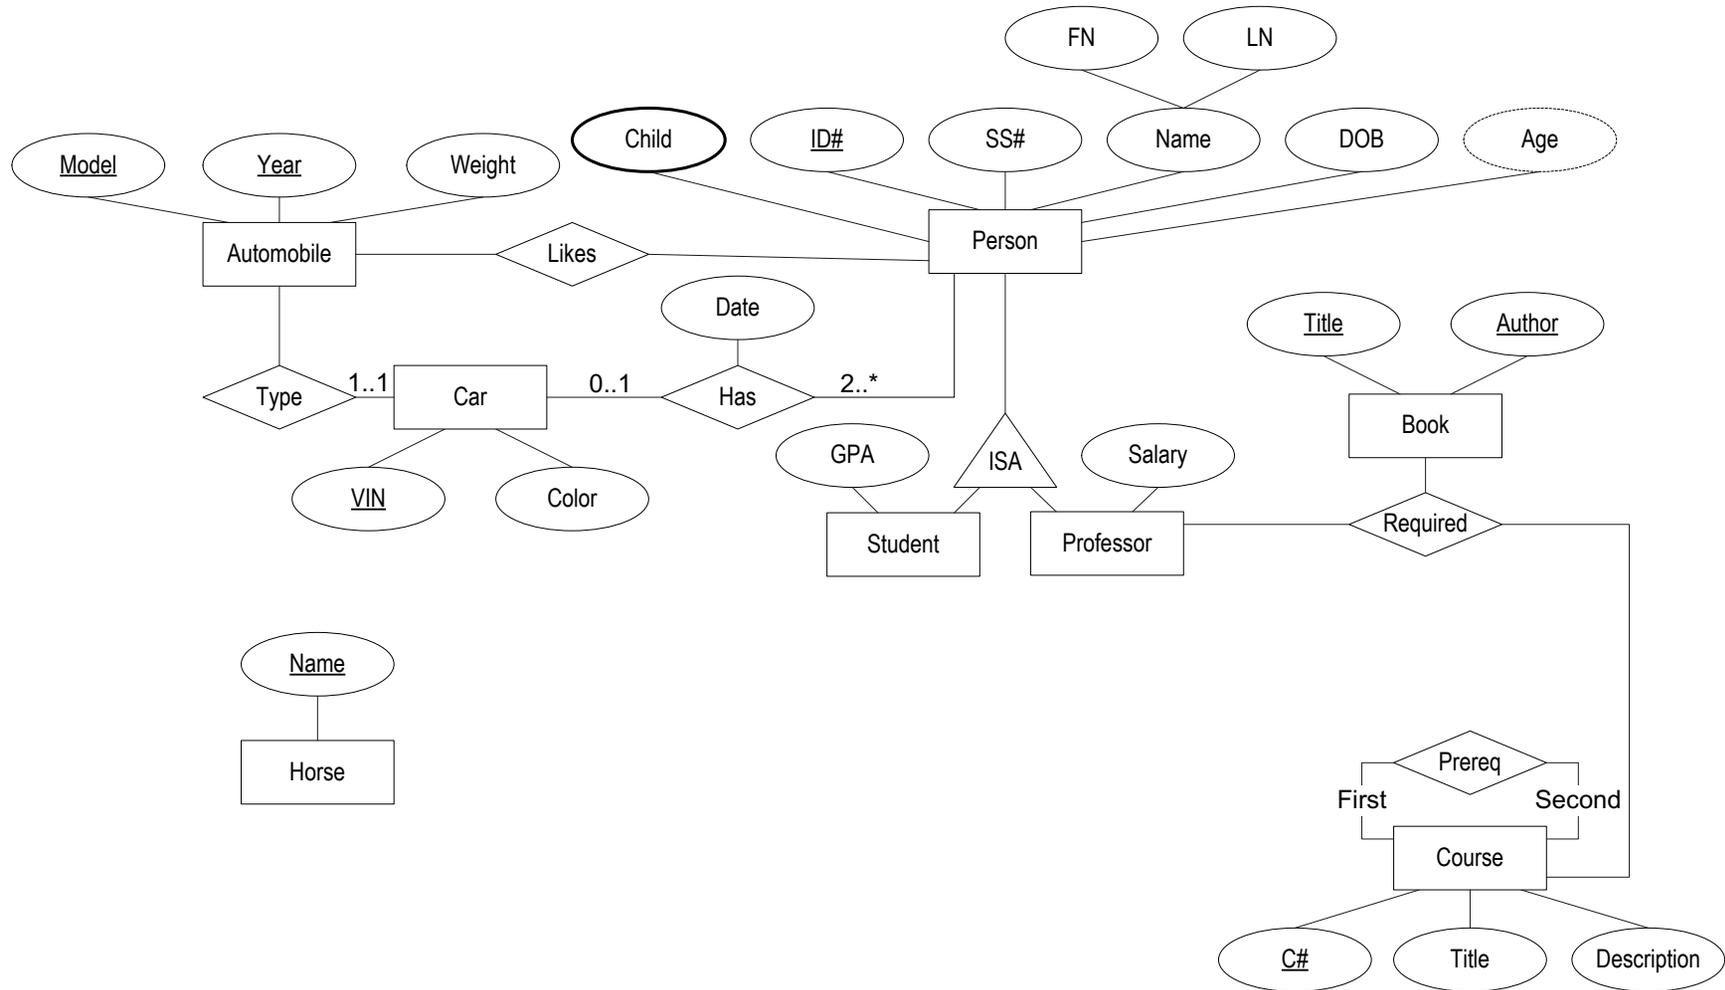
Our ER Diagram



Required

- ◆ **Required**; relationship
 - ◆ Relationship among/between:
 - **Professor**
 - **Course**
 - **Book**
 - ◆ Attributes
 - ◆ Constraints
-
- ◆ A professor specifies that a book is required for a course

Our ER Diagram



Required

- ◆ Note that there are no cardinality or other restrictions
- ◆ Any professor can require any book for any course and a book can be specified by different professors for the same course
- ◆ A book does not have to be required for any course

Section

- ◆ **Section**; entity set
- ◆ Attributes:
 - **Year**
 - **Semester**
 - **Sec#**
 - **MaxSize**
- ◆ Constraints
 - Discriminant: Year, Semester, Sec#
 - Identified through relationship Offered to Course
 - Each Course has to have at least one Section (we have a policy of not putting a course in a catalog unless it has been offered at least once)

Section

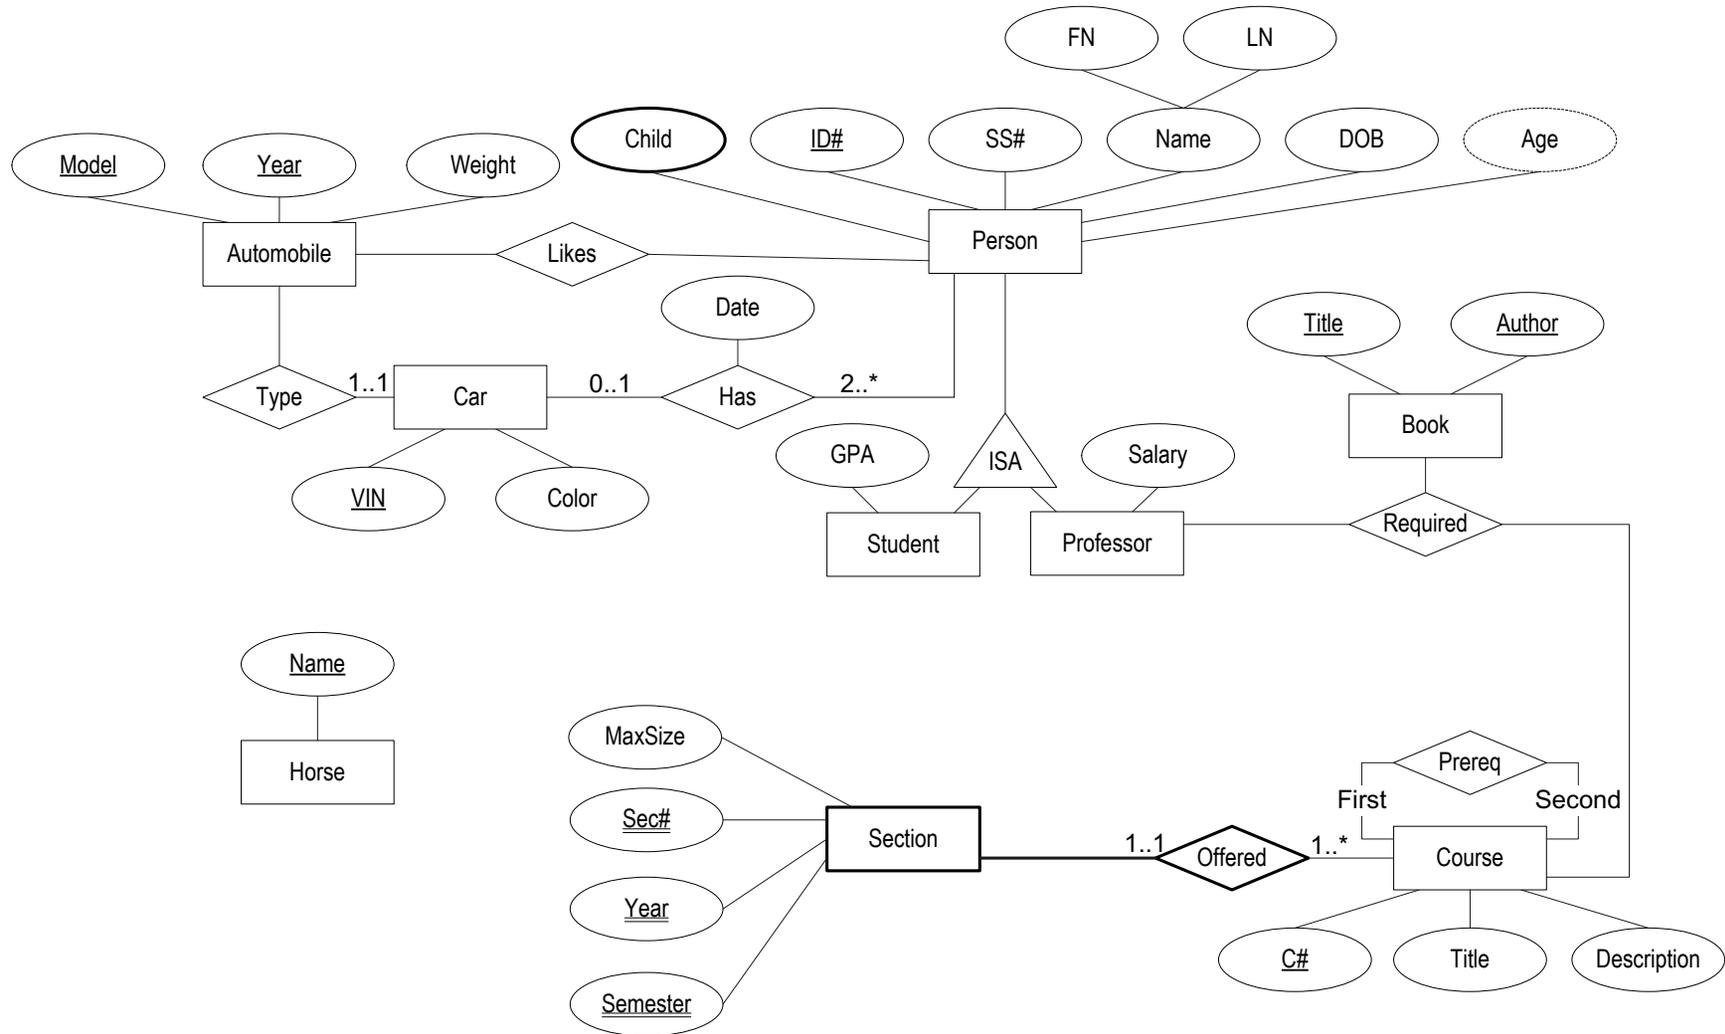
- ◆ Section is a weak entity
- ◆ It is related for the purpose of identification to a strong entity Course by a new relationship Offered
- ◆ It has a discriminant, so it is in fact identified by having the following specified
C#, Year, Semester, Sec#
- ◆ Our current section is identified by:
CSCI-GA.2433, 2013, Fall, 001

Offered

- ◆ **Offered**; relationship
- ◆ Relationship among/between:
 - **Course**
 - **Section**
- ◆ Attributes
- ◆ Constraints
 - Course has to be related to at least one section (see above)
 - Section has to be related to exactly one course (this automatically follows from the fact that section is identified through exactly one course, so maybe we do not need to say this)

- ◆ Note: May be difficult to see, but Section and Offered are both drawn with thick lines

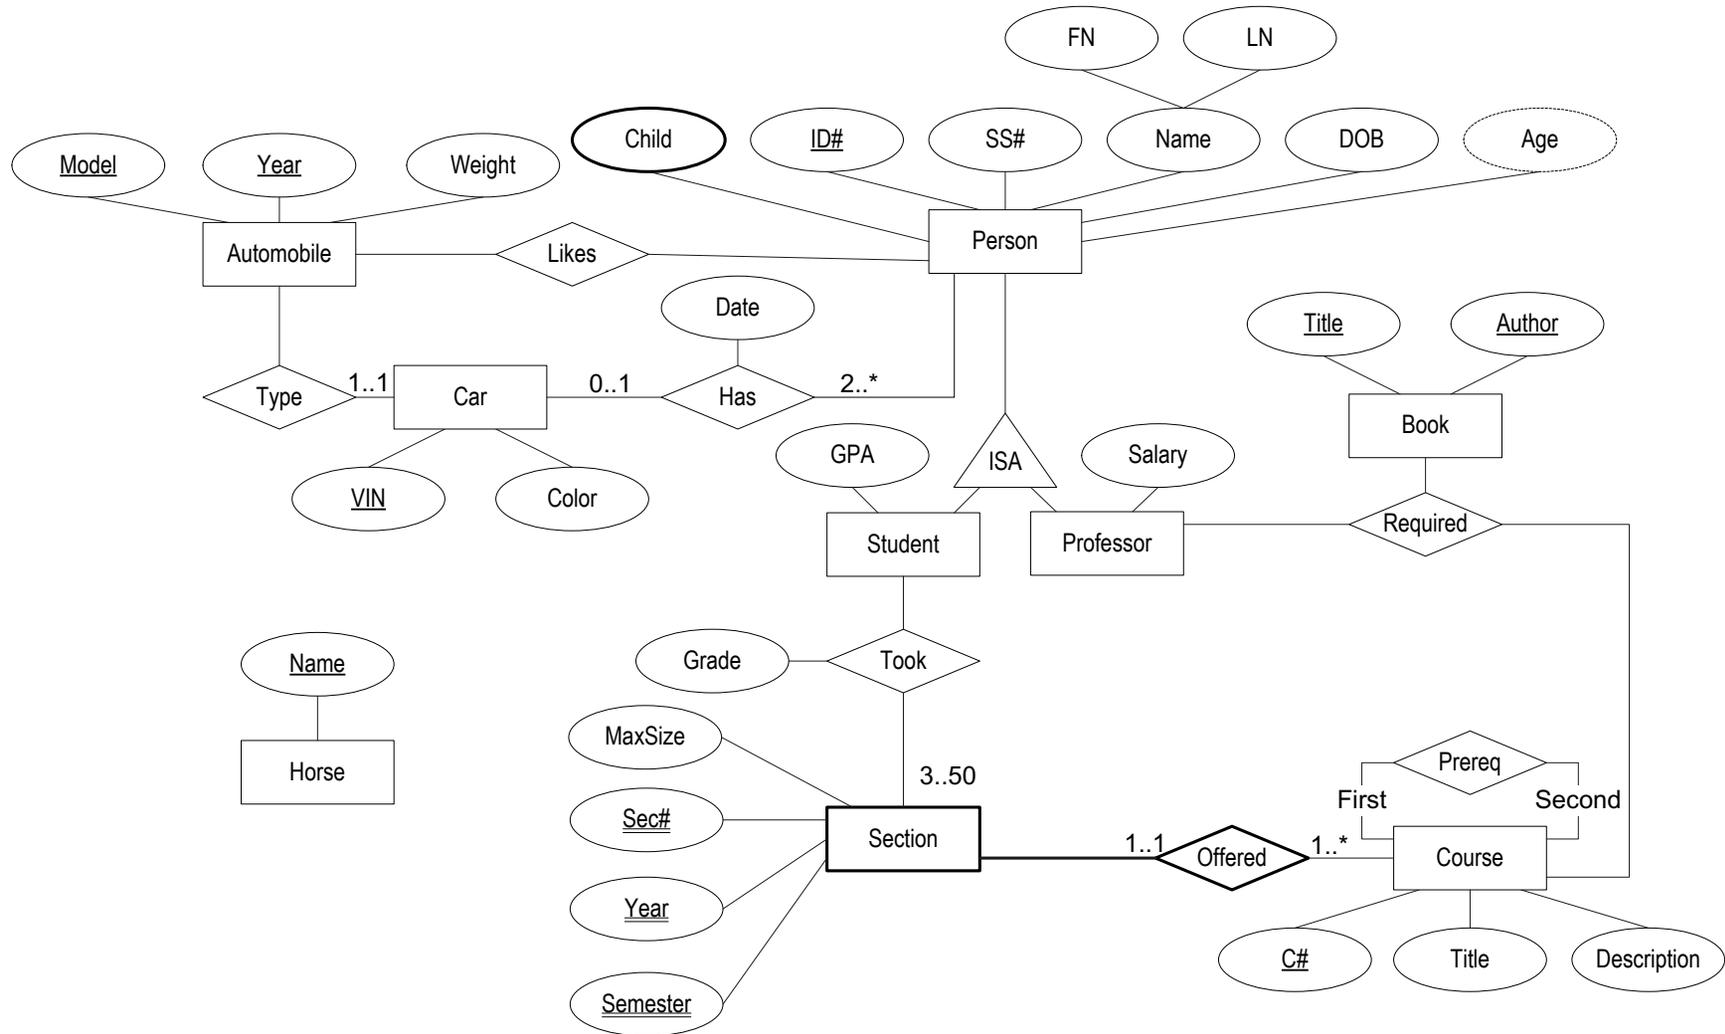
Our ER Diagram



Took

- ◆ **Took**; relationship
- ◆ Relationship among/between
 - **Student**
 - **Section**
- ◆ Attributes
 - **Grade**
- ◆ Constraints
 - Cardinality: 3..50 between Section and Took (this means that a section has between 3 and 50 students)

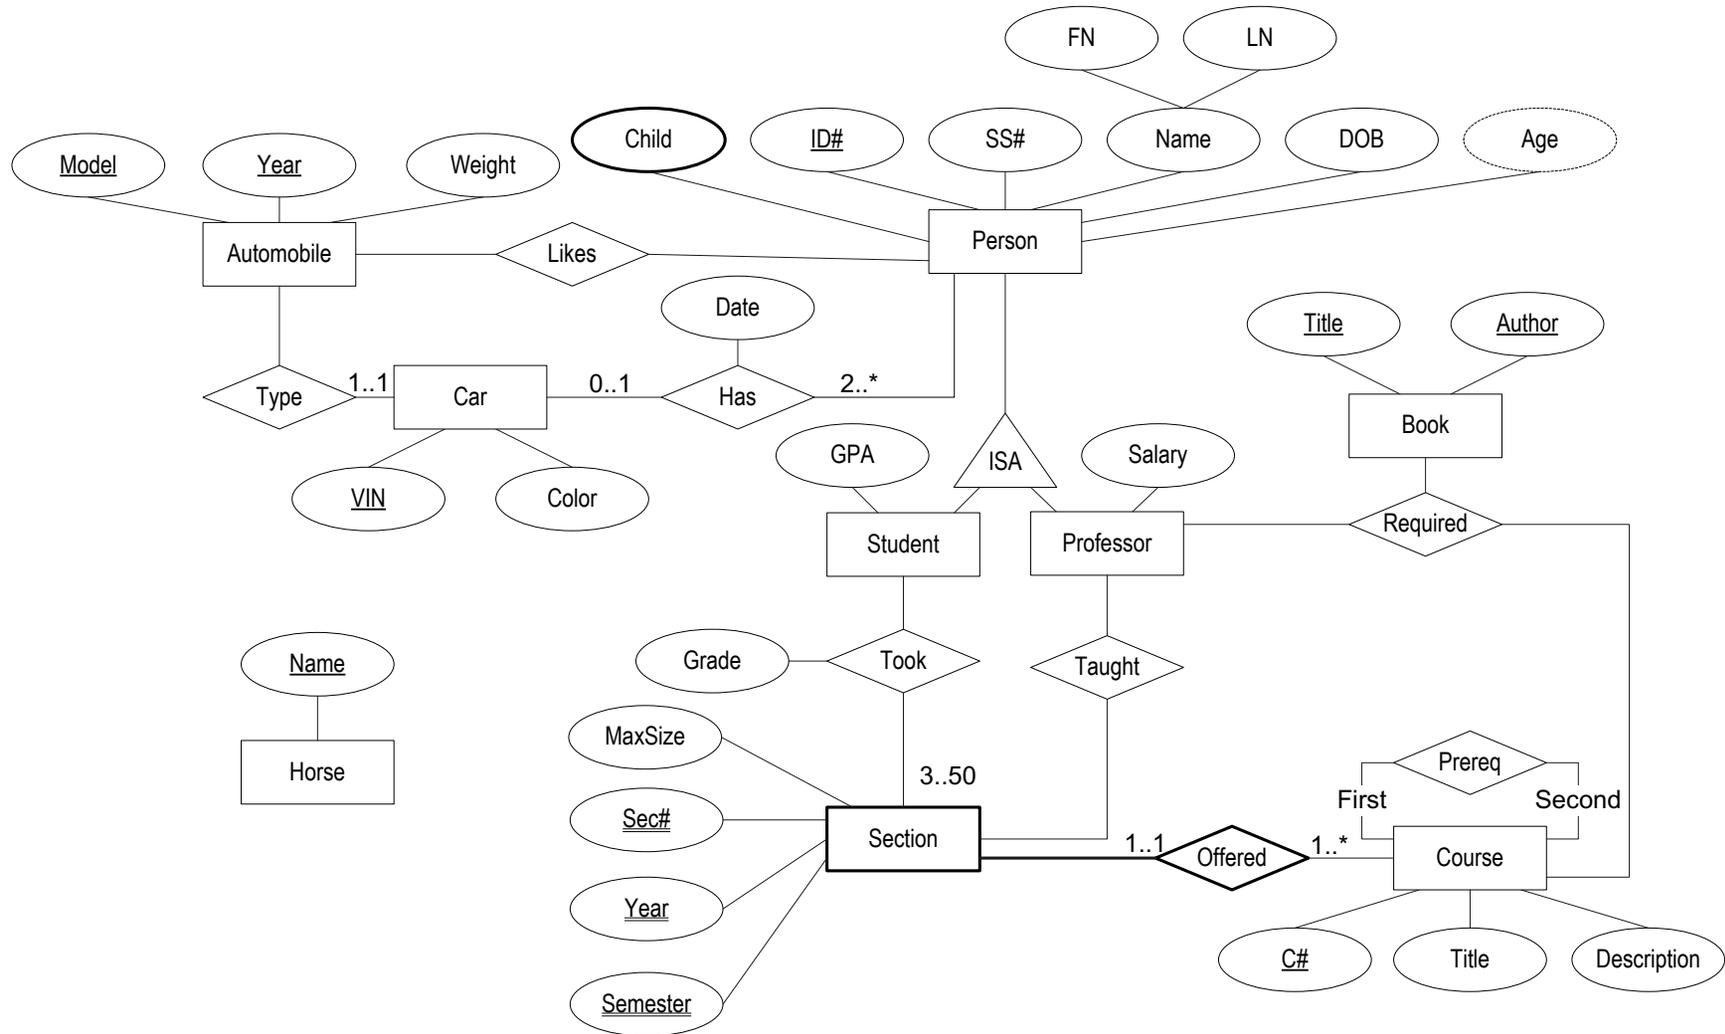
Our ER Diagram



Taught

- ◆ *Taught*; relationship
 - ◆ Relationship among/between
 - *Professor*
 - *Section*
 - ◆ Attributes
-
- ◆ This tells us which professor teach which sections
 - Note there is no cardinality constraint: any number of professors, including zero professors can teach a section (no professor yet assigned, or hypothetical situation)
 - If we wanted, we could have put 1..* between Section and Taught to specify that at least one professor has to be assigned to each section

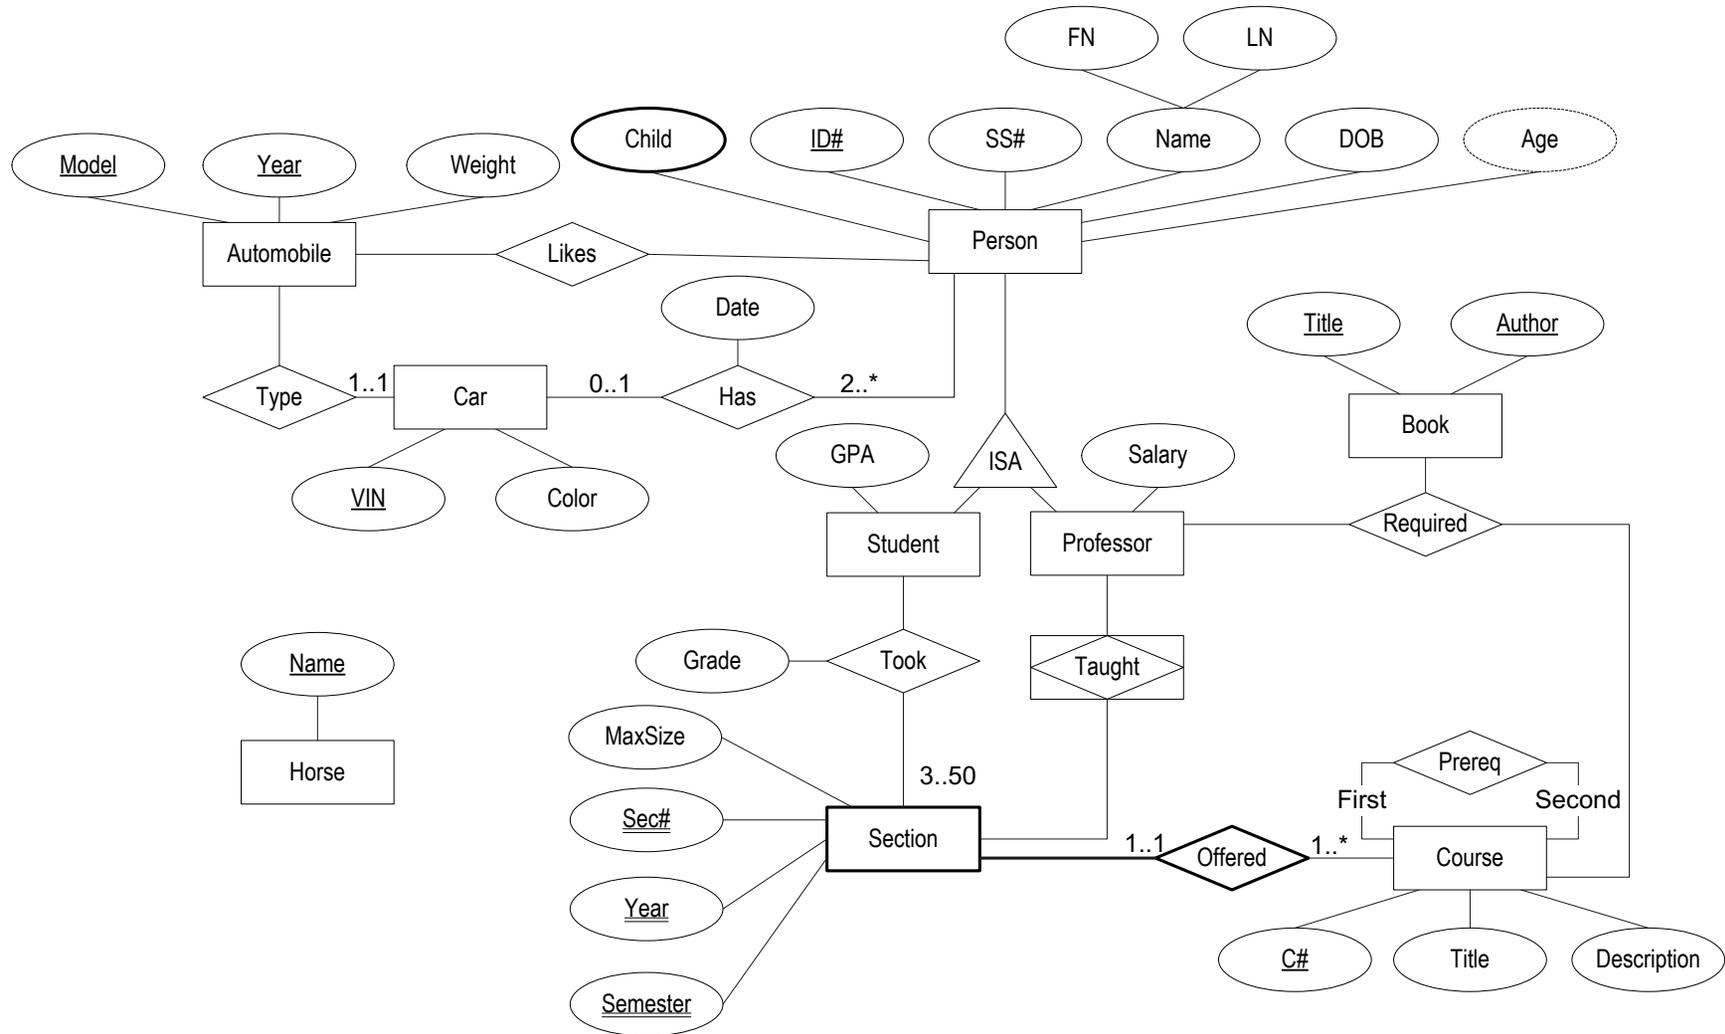
Our ER Diagram



Taught

- ◆ We want to think of Taught as an entity
 - We will see soon why

Our ER Diagram

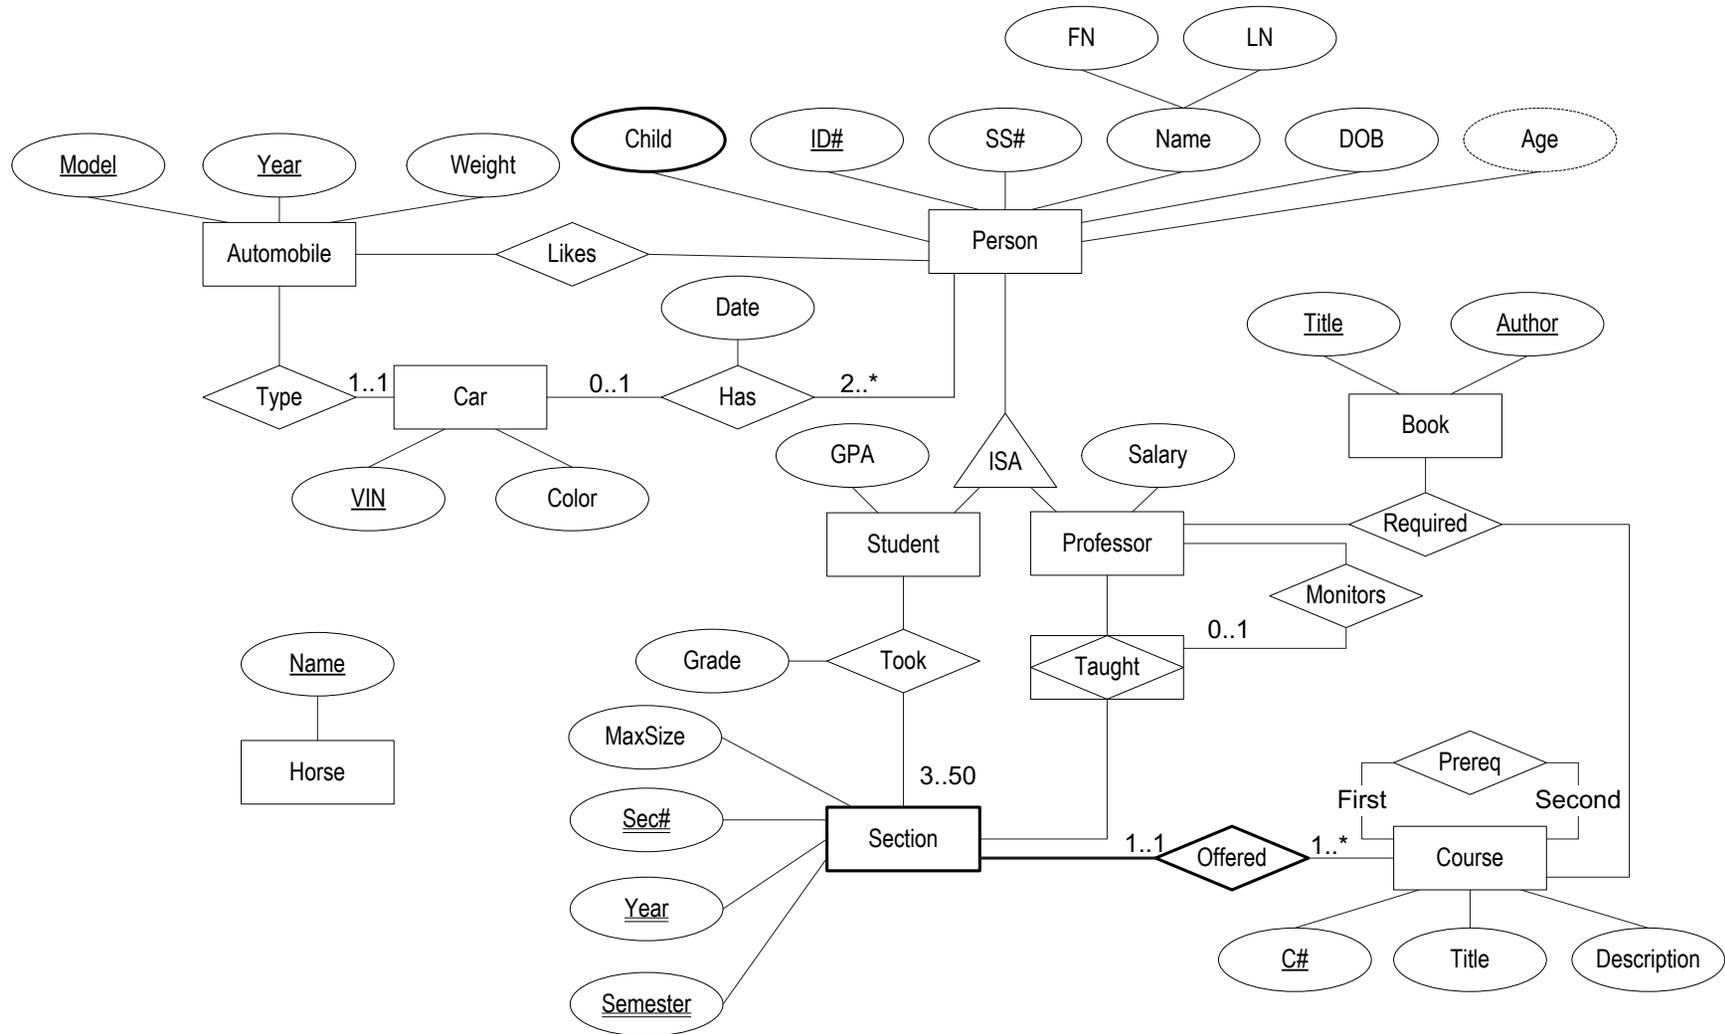


Monitors

- ◆ **Monitors**; relationship
- ◆ Relationship among/between
 - **Professor**
 - **Taught** (considered as an entity)
- ◆ Attributes
- ◆ Constraints
 - Cardinality: 0..1 between Taught and Professor

- ◆ This models the fact that Taught (really a teaching assignment) may need to be monitored by a professor and at most one professor is needed for such monitoring
 - We are not saying whether the professor monitoring the assignment has to be different from the teaching professor in this assignment (but we could do it in SQL DDL, as we shall see later)

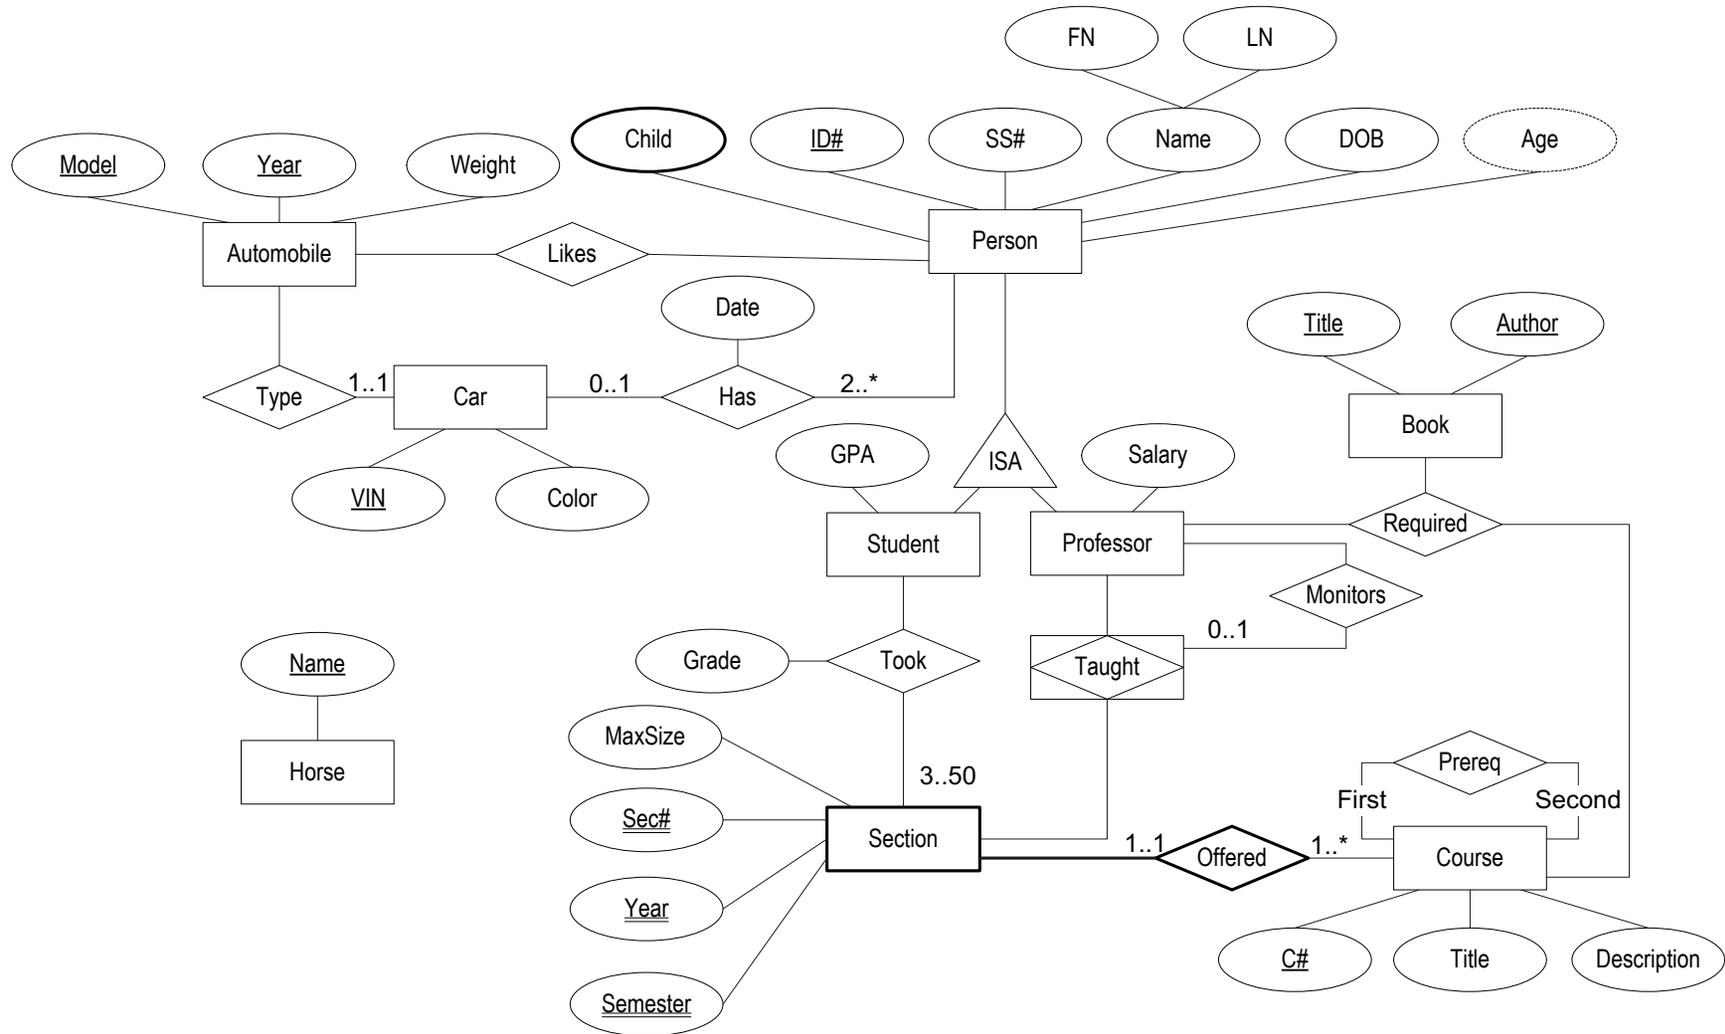
Our ER Diagram



What Can We Learn From The Diagram?

- ◆ Let's look
- ◆ We will review everything we can learn just by looking at the diagram

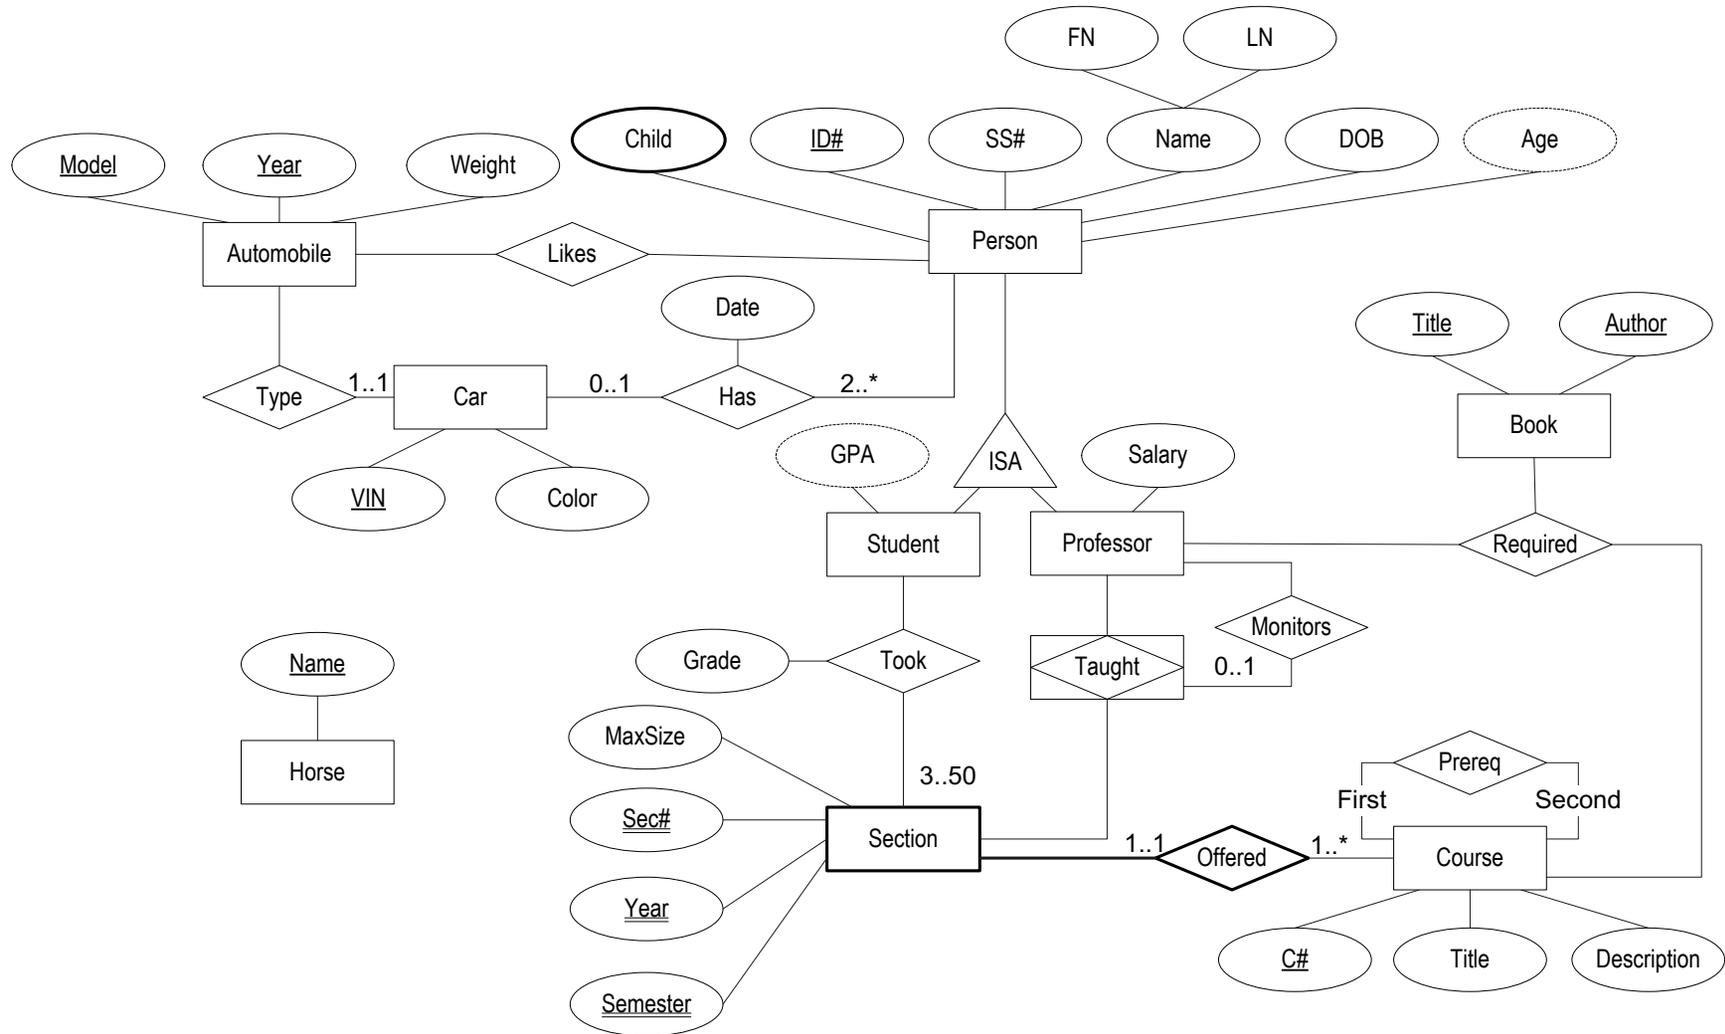
Our ER Diagram



GPA

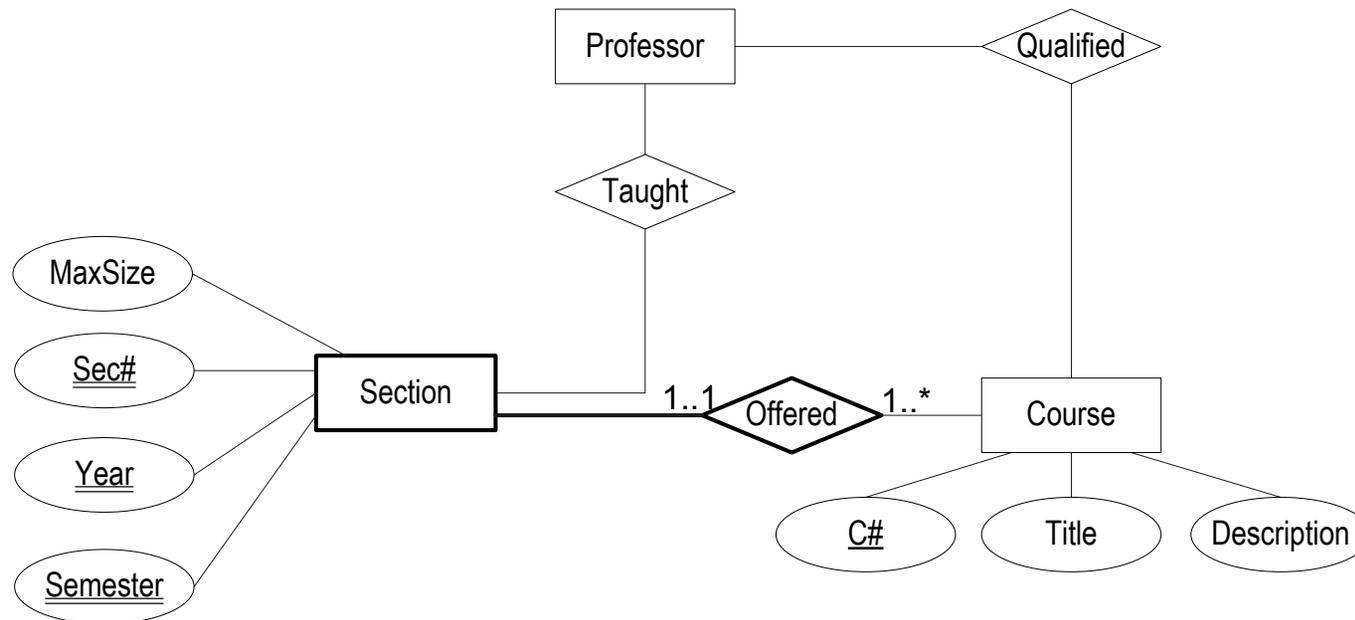
- ◆ We now observe that GPA should probably be modeled as a derived attribute, as it is computed from the student's grade history
- ◆ So, we may want to revise the diagram

Our ER Diagram



Some Constraints Are Difficult To Specify

- ◆ Imagine that we also have relationship Qualified between Professor and Course specifying which professors are qualified to teach which courses
- ◆ We probably use words and not diagrams to say that only a qualified professor can teach a course



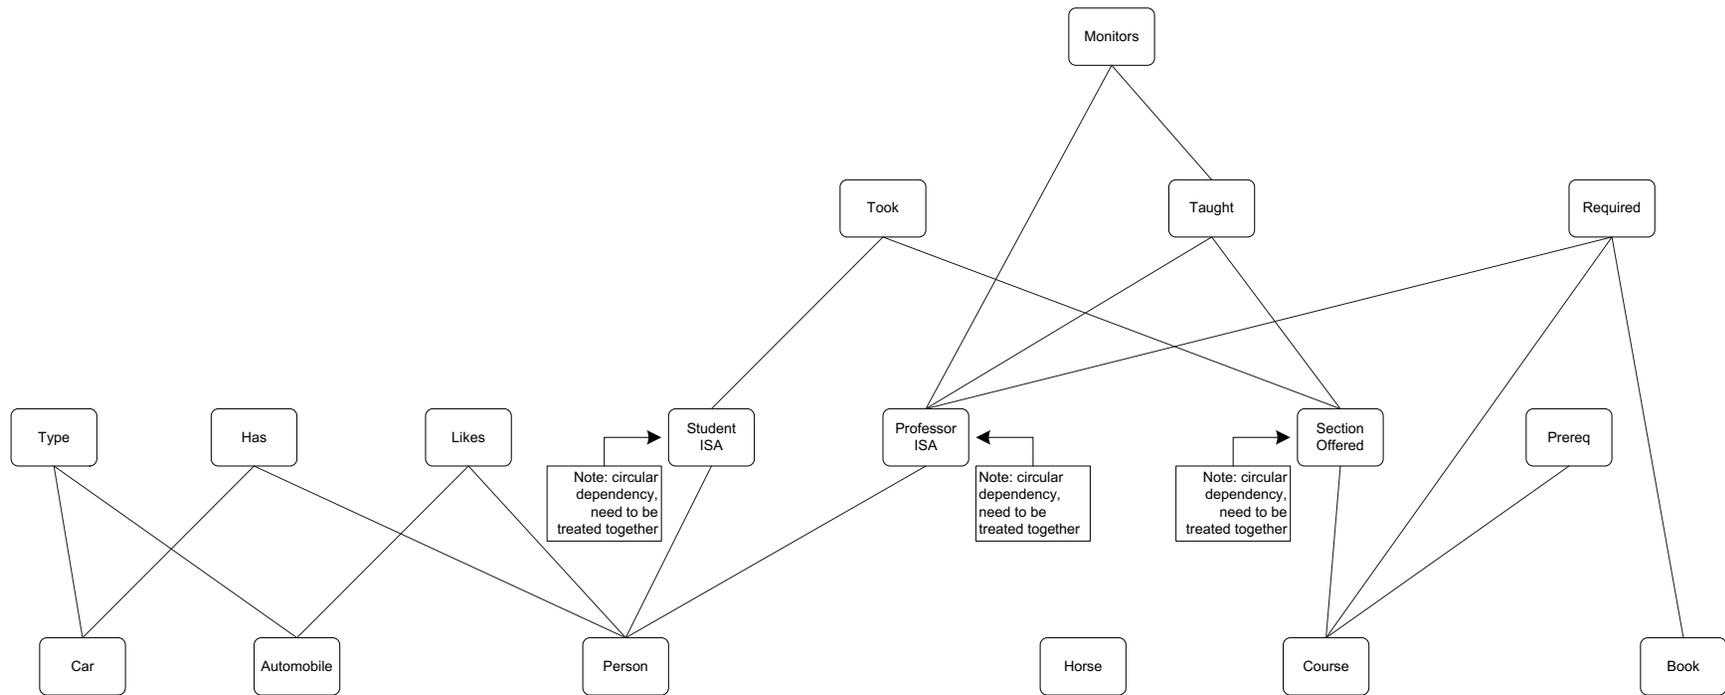
Annotate, Annotate, Annotate ...

- ◆ **An ER diagram should be annotated with all known constraints**

Hierarchy For Our ER Diagram

- ◆ There is a natural hierarchy for our ER diagram
- ◆ It shows us going from bottom to top how the ER diagram was constructed
- ◆ Section and Offered have to be constructed together as there is a circular dependency between them
- ◆ Similar issue comes up when dealing with ISA

Hierarchy For Our ER Diagram



Next

- ◆ We will learn how to take an ER diagram and convert it into a relational database
- ◆ We will learn how to specify such databases using Visio (which you will get for free from NYU)

Key Ideas

- ◆ ER diagrams
- ◆ Entity and Entity Set
- ◆ Attribute
 - Base
 - Derived
 - Simple
 - Composite
 - Singlevalued
 - Multivalued
- ◆ Superkey
- ◆ Key
- ◆ Candidate Key
- ◆ Primary Key
- ◆ UNIQUE

Key Ideas

- ◆ Relationship
- ◆ Binary relationship and its functionality
- ◆ Non-binary relationship
- ◆ Relationship with attributes
- ◆ Aggregation
- ◆ Strong and weak entities
- ◆ Discriminant
- ◆ ISA
 - Disjoint
 - Overlapping
 - Total
 - Partial
 - Specialization
 - Generalization

Key Ideas

- ◆ General Cardinality Constraints
- ◆ Case study of modeling