

Pydantic

Pydantic makes type hints validate your data at runtime, covered here from coercion and constraints to validators and serialization.

Declare a BaseModel

Subclass BaseModel; your type hints become the schema.

PURPOSE	COMMAND	RESULT
Define a model from type hints.	<pre>class User(BaseModel): id:int name:str</pre>	
Add a default value.	<pre>name:str = "anon"</pre>	
Make a field nullable.	<pre>nickname:str None = None</pre>	
Construct an instance.	<pre>u =User(id=1, name="ada")</pre>	
List the declared fields.	<pre>User.model_fields</pre>	
Access attributes (typed).	<pre>u.id # int u.name # str</pre>	

each annotated attribute becomes a validated field; v2 replaces the inner class Config with model_config = ConfigDict(...)

Validate Input

model_validate parses untrusted dicts and JSON into typed objects.

PURPOSE	COMMAND	RESULT
Validate a Python dict.	<pre>User.model_validate({"id": 1, ...})</pre>	
Validate a JSON string.	<pre>User.model_validate_json(s)</pre>	
Construct from keywords.	<pre>User(id=1, name="ada")</pre>	
Validate a list of items.	<pre>TypeAdapter(list[User]).validate_python(rows)</pre>	
Bad input never passes silently.	<pre>User.model_validate({"id": "oops"})</pre>	
Forbid coercion (strict).	<pre>model_validate({"id": "1"},strict=True)</pre>	

model_validate (dict) and model_validate_json (text) replace v1 parse_obj / parse_raw; TypeAdapter validates non-model types

Coerce & Constrain

Field() adds bounds, defaults, and metadata to a type.

PURPOSE	COMMAND	RESULT
Coerce a string to the type.	<code>age: int... Account(age="30")</code>	
Bound a numeric field.	<code>age: int =Field(ge=0, le=120)</code>	
Constrain string length.	<code>name: str =Field(min_length=1, max_length=50)</code>	
Safe mutable default.	<code>tags: list[str] =Field(default_factory=list)</code>	
Regex via Annotated.	<code>Annotated[str, Field(pattern=r"^\d{5}\$")]</code>	
Map external key names.	<code>Field(validation_alias=AliasChoices("user_id", "userId"))</code>	

pydantic coerces sensible values by default; Field() layers ge/le, min_length/max_length, pattern, aliases, and safe defaults

Read a ValidationError







One exception lists every problem with loc, msg, and type.

PURPOSE	COMMAND	RESULT												
Catch the failure.	<code>except ValidationErrors as e:</code>													
Count the problems.	<code>e.error_count()</code>													
Read structured error rows.	<code>e.errors()</code>	<table border="1"> <thead> <tr> <th>loc</th> <th>type</th> <th>msg</th> <th>input</th> </tr> </thead> <tbody> <tr> <td>age</td> <td>less than</td> <td>--</td> <td>200</td> </tr> <tr> <td>name</td> <td>missing</td> <td>--</td> <td>--</td> </tr> </tbody> </table>	loc	type	msg	input	age	less than	--	200	name	missing	--	--
loc	type	msg	input											
age	less than	--	200											
name	missing	--	--											
Pinpoint the bad field.	<code>e.errors()[0]["loc"]</code>													
Human-readable summary.	<code>print(e)</code>													
JSON-shaped error report.	<code>e.json()</code>													

pydantic collects every problem into one ValidationError; each error dict carries loc, type, msg, and input

Field & Model Validators


Hook custom rules per field or across the whole model.

PURPOSE	COMMAND	RESULT
Normalize one field.	<code>@field_validator("name") → strip().lower()</code>	
Validate after coercion.	<code>@field_validator("age", mode="after")</code>	
Cross-field rule.	<code>@model_validator(mode="after") → self</code>	
Pre-process raw input.	<code>@model_validator(mode="before")</code>	
Derive a read-only field.	<code>@computed_field over a @property</code>	
Raise to reject a value.	<code>raise ValueError("bad")</code>	

@field_validator (wrap a classmethod) runs per field; @model_validator runs across the instance; raise ValueError to reject

Nest & Reuse Models







Models compose: nest them, list them, union them.

PURPOSE	COMMAND	RESULT
Nest one model in another.	<code>addresses: list[Address]</code>	
Validate nested dicts deeply.	<code>Person.model_validate({"addresses": [...]})</code>	
Accept one of several types.	<code>value: int str</code>	
Tagged (discriminated) union.	<code>Field(discriminator="kind") Cat Dog</code>	
Reuse fields via inheritance.	<code>class Admin(User): level: int</code>	
Recursive / self-reference.	<code>children: list["Node"] → model_rebuild()</code>	

nested dicts validate all the way down into real sub-objects; a discriminator routes tagged unions without guessing

Serialize Out





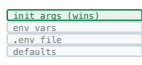

`model_dump` for a dict, `model_dump_json` for a JSON string.

PURPOSE	COMMAND	RESULT
Dump to a Python dict.	<code>u.model_dump()</code>	 User → { } real types
Dump to a JSON string.	<code>u.model_dump_json()</code>	 User → "..." dates → ISO
JSON-safe values in a dict.	<code>u.model_dump(mode="json")</code>	 {"when": "2026-06-18T09:00"}
Drop empty / unset fields.	<code>u.model_dump(exclude_none=True)</code>	 id: 1, name: None → slimer
Hide a field from output.	<code>secret: str = Field(exclude=True)</code>	 id, name, secret → never serialized
Emit the JSON Schema.	<code>User.model_json_schema()</code>	 {"properties": ..., "required": [...]} OpenAPI

`model_dump` keeps rich Python types; `model_dump_json` renders them as strings; `mode="json"` gives JSON-safe values in a dict

Config with pydantic-settings

`BaseSettings` loads typed config from env vars and `.env` files.

PURPOSE	COMMAND	RESULT
Install the separate package.	<code>uv add pydantic-settings</code>	 its own package (split from core in v2)
Declare a typed settings class.	<pre>class Settings(BaseSettings): port: int = 8000</pre>	 Settings, port: int = 8000, env-aware
Pull values from the environment.	<code>Settings() ← APP_PORT=9000</code>	 APP_PORT=9000 → Settings → port=9000
Namespace and load a <code>.env</code> file.	<code>SettingsConfigDict(env_prefix="APP_", env_file=".env")</code>	 .env, APP → SettingsConfigDict, file + prefix
Source precedence.	<code>init args > env vars > .env > defaults</code>	 init args (wins), env vars, .env file, defaults
Validate config like any model.	<code>APP_PORT=oops → ValidationError</code>	 PORT=oops → ValidationError, int_parsing fail fast

`BaseSettings` (a separate package in v2) reads typed config from env and `.env`, coercing and validating it like any `BaseModel`