Trivial proof of knowledge: prover sends witness in the clear to the verifier
↳ In most applications, we additionally require zero-knowledge

Note: knowledge is a strictly stronger property than soundness
↳ if protocol has knowledge error $\varepsilon \Rightarrow$ it also has soundness error ε (i.e. a dishonest prover convinces an honest verifier of a false statement with probability at most ε)

Proving knowledge of discrete log (Schnorr's protocol)

assume $g, h \in \mathbb{G}$ where $\mathbb{G}$ has prime order q

Suppose prover wants to prove it knows x such that $h = g^x$ (i.e. prover demonstrates knowledge of discrete log of h base g)

Prover		Verifier
$r \xleftarrow{R} \mathbb{Z}_p$, $u \leftarrow g^r$	$\xrightarrow{u}$	
	$\xleftarrow{c}$	$c \xleftarrow{R} \mathbb{Z}_p$
$z \leftarrow r + cx$	$\xrightarrow{z}$	

verify that $g^z = u \cdot h^c$

Completeness: if $z = r + cx$, then
$$g^z = g^{r+cx} = g^r g^{cx} = u \cdot h^c$$

zero knowledge only required to hold against an honest verifier (e.g., view of the honest verifier can be simulated)

Honest-Verifier Zero-Knowledge: build a simulator as follows (familiar strategy: run the protocol in "reverse"):

on input (g, h):
1. Sample $z \xleftarrow{R} \mathbb{Z}_p$
2. Sample $c \xleftarrow{R} \mathbb{Z}_p$
3. set $u = g^z / h^c$ and output (u, c, z)

uniformly random challenge

uniformly random group element since z is uniformly random

chosen so that $g^z = u \cdot h^c$ (relation satisfied by a valid proof)

simulated transcript is identically distributed as the real transcript with an honest verifier

What goes wrong if the challenge is not sampled uniformly at random (i.e., if the verifier is dishonest)
Above simulation no longer works (since we cannot sample z first)
↳ To get general zero knowledge, we require that the verifier first commit to its challenge (using a statistically hiding commitment)

Knowledge: Suppose P^* is (possibly malicious) prover that convinces honest verifier with probability 1. We construct an extractor as follows:

for simplicity, we assume P^* succeeds with probability 1

1. Run the prover P^* to obtain an initial message u.
2. Send a challenge $c_1 \xleftarrow{R} \mathbb{Z}_p$ to P^*. The prover replies with a response z_1.
3. "Rewind" the prover P^* so its internal state is the same as it was at the end of Step 1. Then, send another challenge $c_2 \xleftarrow{R} \mathbb{Z}_p$ to P^*. Let z_2 be the response of P^*.
4. Compute and output $x = (z_1 - z_2)(c_1 - c_2)^{-1} \in \mathbb{Z}_p$.

Since P^* succeeds with probability 1 and the extractor perfectly simulates the honest verifier's behavior, with probability 1, both (u, c_1, z_1) and (u, c_2, z_2) are both accepting transcripts. This means that

$$g^{z_1} = u \cdot h^{c_1} \quad \text{and} \quad g^{z_2} = u \cdot h^{c_2}$$

$$\Longrightarrow \frac{g^{z_1}}{h^{c_1}} = \frac{g^{z_2}}{h^{c_2}} \Longrightarrow g^{z_1 + c_2 x} = g^{z_2 + c_1 x}$$

with overwhelming probability, $c_1 \neq c_2$

$$\Longrightarrow x = (z_1 - z_2)(c_1 - c_2)^{-1} \in \mathbb{Z}_p$$

Thus, extractor succeeds with overwhelming probability.

(Boneh-Shoup, Lemma 19.2)

If P^* succeeds with probability ε, then need to rely on "Rewinding Lemma" to argue that extractor obtains two accepting transcripts with probability at least $\varepsilon^2 - 1/p$.

How can a prover both prove knowledge and yet be zero-knowledge at the same time?
- ↳ Extractor operates by "rewinding" the prover (if the prover has good success probability, it can answer most challenges correctly.
- ↳ But in the real (actual) protocol, verifier cannot rewind (i.e., verifier only sees prover on fresh protocol executions), which can provide zero-knowledge.

Identification protocol from discrete log:

client (x) — client's secret (credential)

server ($g, h = g^x$) — public verification key

protocol is precisely 3-round Schnorr proof of knowledge of discrete log

Essentially, the discrete log of h (base g) is the client's "password" and instead of sending the password in the clear to the server, the client proves in zero-knowledge that it knows x

Correctness of this protocol follows from completeness of Schnorr's protocol

(Active) security follows from knowledge property and zero-knowledge
- ↳ Intuitively: knowledge says that any client that successfully authenticates must know secret x

  zero-knowledge says that interactions with honest client (i.e., the prover) do not reveal anything about x

  (for active security, require protocol that provides general zero-knowledge rather than just HVZK)